

THE COMPLEXITY OF GLOBAL CARDINALITY CONSTRAINTS

ANDREI A. BULATOV^a AND DÁNIEL MARX^b

^a School of Computing Science, Simon Fraser University, Burnaby, Canada
e-mail address: abulatov@cs.sfu.ca

^b Tel Aviv University, Tel Aviv, Israel
e-mail address: dmarx@cs.bme.hu

ABSTRACT. In a constraint satisfaction problem (CSP) the goal is to find an assignment of a given set of variables subject to specified constraints. A global cardinality constraint is an additional requirement that prescribes how many variables must be assigned a certain value. We study the complexity of the problem CCSP(Γ), the constraint satisfaction problem with global cardinality constraints that allows only relations from the set Γ . The main result of this paper characterizes sets Γ that give rise to problems solvable in polynomial time, and states that the remaining such problems are NP-complete.

1. INTRODUCTION

In a constraint satisfaction problem (CSP) we are given a set of variables, and the goal is to find an assignment of the variables subject to specified constraints. A constraint is usually expressed as a requirement that combinations of values of a certain (usually small) set of variables belong to a certain relation. CSPs have been intensively studied in both theoretical and practical perspectives. On the theoretical side the key research direction has been the complexity of the CSP when either the interaction of sets constraints are imposed on, that is, the hypergraph formed by these sets, is restricted [16, 17, 18], or restrictions are on the type of allowed relations [21, 9, 5, 6, 2]. In the latter direction the main focus has been on the so called *Dichotomy conjecture* [14] suggesting that every CSP restricted in this way is either solvable in polynomial time or is NP-complete.

This ‘pure’ constraint satisfaction problem is sometimes not enough to model practical problems, as some constraint that have to be satisfied are not ‘local’ in the sense that they cannot be viewed as applied to only a limited number of variables. Constraints of this type are called *global*. Global constraints are very diverse, the current Global Constraint Catalog (see <http://www.emn.fr/x-info/sdemasse/gccat/>) lists 348 types of such constraints. In this paper we focus on *global cardinality constraints* [24, 4, 25]. A global cardinality constraint π is specified for a set of values D and a set of variables V , and is given by a mapping $\pi : D \rightarrow \mathbb{N}$ that assigns a natural number to each element of D such that $\sum_{a \in D} \pi(a) = |V|$.

1998 ACM Subject Classification: F.2.2, F.4.1.

2000 Mathematics Subject Classification: 68Q25, 68W40.

Key words and phrases: constraint satisfaction problem, cardinality constraints, complexity.

An assignment of variables V satisfies π if for each $a \in D$ the number of variables that take value a equals $\pi(a)$. In a CSP with global cardinality constraints, given a CSP instance and a global cardinality constraint π , the goal is to decide if there is a solution of the CSP instance satisfying π . The restricted class of CSPs with global cardinality constraints such that every instance from this class uses only relations from a fixed set Γ of relations (such a set is often called a *constraint language*) is denoted by $\text{CCSP}(\Gamma)$. We consider the following problem: Characterize constraint languages Γ such that $\text{CCSP}(\Gamma)$ is solvable in polynomial time. More general versions of global cardinality constraints have appeared in the literature, see, e.g. [24], where the number of variables taking value a has to belong to a prescribed set of cardinalities (rather than being exactly $\pi(a)$). In this paper we call the CSP allowing such generalized constraints *extended CSP with cardinality constraints*. As we discuss later, our results apply to this problem as well.

The complexity of $\text{CCSP}(\Gamma)$ has been studied in [12] for constraint languages Γ on a 2-element set. It was shown that $\text{CCSP}(\Gamma)$ is solvable in polynomial time if every relation in Γ is *width-2-affine*, i.e. it can be expressed as the set of solutions of a system of linear equations over a 2-element field containing at most 2 variables, or, equivalently, using the equality and disequality clauses; otherwise it is NP-complete (we assume $P \neq \text{NP}$). In the 2-element case $\text{CCSP}(\Gamma)$ is also known as the k -ONES(Γ) problem, since a global cardinality constraint can be expressed by specifying how many ones (the set of values is thought to be $\{0, 1\}$) one wants to have among the values of variables. The parameterized complexity of k -ONES(Γ) has also been studied [23], where k is used as a parameter.

In the case of a 2-element domain, the polynomial classes can be handled by a standard application of dynamic programming. Suppose that the instance is given by a set of unary clauses and binary equality/disequality clauses. Consider the graph formed by the binary clauses. There are at most two possible assignments for each connected component of the graph: setting the value of a variable uniquely determines the values of all the other variables in the component. Thus the problem is to select one of the two assignments for each component. Trying all possibilities would be exponential in the number of components. Instead, for $i = 1, 2, \dots$, we compute the set S_i of all possible pairs (x, y) such that there is a partial solution on the first i components containing exactly x zeroes and exactly y ones. It is not difficult to see that S_{i+1} can be computed if S_i is already known.

We generalize the results of [12] for arbitrary finite sets and arbitrary constraint languages. As usual, the characterization for arbitrary finite sets is significantly more complex and technical than for the 2-element set. As a straightforward generalization of the 2-element case, we can observe that the problem is polynomial-time solvable if every relation can be expressed by binary mappings. In this case, setting a single value in a component uniquely determines all the values in the component. Therefore, if the domain is D , then there are at most $|D|$ possible assignments in each component, and the same dynamic programming technique can be applied (but this time the set S_i contains $|D|$ -tuples instead of pairs).

One might be tempted to guess that the class described in the previous paragraph is the only class where CCSP is polynomial-time solvable. However, it turns out that there are more general tractable classes. First, suppose that the domain is partitioned into equivalence classes, and the binary constraints are mappings between the equivalence classes. This means that the values in the same equivalence class are completely interchangeable. Thus it is sufficient to keep one representative from each class, and then the problem can be solved by the algorithm sketched in the previous paragraph. Again, one might believe that

this construction gives all the tractable classes, but the following example shows that it does not cover all the tractable cases.

Example 1.1. Let $R = \{(1, 2, 3), (1, 4, 5), (a, b, c), (d, e, c)\}$. We claim that $\text{CCSP}(\{R\})$ is polynomial-time solvable. Consider the graph on the variables where two variables are connected if and only if they appear together in a constraint. As before, for each component, we compute a set containing all possible cardinality vectors, and then use dynamic programming. In each component, we have to consider only two cases: either every variable is in $\{1, 2, 3, 4, 5\}$ or every variable is in $\{a, b, c, d, e\}$. If every variable of component K is in $\{1, 2, 3, 4, 5\}$, then R can be expressed by the unary constant relation $\{1\}$, and the binary relation $R' = \{(2, 3), (4, 5)\}$. The binary relations partition component K into sub-components K_1, \dots, K_t . Since R' is a mapping, there are at most 2 possible assignments for each sub-component. Thus we can use dynamic programming to compute the set of all possible cardinality vectors on K that use only the values in $\{1, 2, 3, 4, 5\}$. If every variable of K is in $\{a, b, c, d, e\}$, then R can be expressed as the unary constant relation $\{c\}$ and the binary relation $R'' = \{(a, b), (d, e)\}$. Again, binary relation R'' partitions K into sub-components, and we can use dynamic programming on them. Observe that the sub-components formed by R' and the sub-components formed by R'' can be different: in the first case, u and v are adjacent if they appear in the second and third coordinates of a constraint, while in the second case, u and v are adjacent if they appear in the first and second coordinates of a constraint.

It is not difficult to make Example 1.1 more complicated in such a way that we have to look at sub-subcomponents and perform multiple levels of dynamic programming. This suggests that it would be difficult to characterize the tractable relations in a simple combinatorial way.

We give two characterizations of finite CCSP, one more along the line of the usual approach to the CSP, using polymorphisms, and another more combinatorial one. The latter is more technical, but it is much more suitable for algorithms.

A polymorphism of a constraint language is an operation that preserves every relation from the language. The types of polymorphisms we need here are quite common and have appeared in the literature many times. A ternary operation m satisfying the equations $m(x, x, y) = m(x, y, x) = m(y, x, x) = x$ is said to be *majority*, and a ternary operation h satisfying $h(x, y, y) = h(y, y, x) = x$ is said to be *Mal'tsev*. An operation is *conservative* if it always takes a value equal to one (not necessarily the same one) of its arguments.

Theorem 1.2. For a constraint language Γ , the problem $\text{CCSP}(\Gamma)$ is polynomial time solvable if and only if Γ has a majority polymorphism and a conservative Mal'tsev polymorphism. Otherwise it is NP-complete.

Observe that for constraint languages over a 2-element domain, Theorem 1.2 implies the characterization of Creignou et al. [12]. Width-2 affine is equivalent to affine and bijunctive (definable in 2SAT), and over a 2-element domain, affine is equivalent to having a conservative Mal'tsev polymorphism and bijunctive is equivalent to having a majority polymorphism.

The second characterization uses logical definability. The right generalization of mappings is given by the notion of rectangularity. A binary relation R is called *rectangular* if $(a, c), (a, d), (b, d) \in R$ implies $(b, c) \in R$. We say that a pair of equivalence relations α and β over the same domain *cross*, if there is an α -class C and a β -class D such that none of

$C \setminus D$, $C \cap D$, and $D \setminus C$ is empty. A relation is *2-decomposable* if it can be expressed as a conjunction of binary relations. We denote by $\langle\langle\Gamma\rangle\rangle$ the set of all relations that are primitive positive definable in Γ . A constraint language is said to be *non-crossing decomposable* if every relation from $\langle\langle\Gamma\rangle\rangle$ is 2-decomposable, every binary relation from $\langle\langle\Gamma\rangle\rangle$ is rectangular, and no pair of equivalence relations from $\langle\langle\Gamma\rangle\rangle$ cross. For detailed definitions and discussion see Section 2.

Theorem 1.3. For a constraint language Γ , the problem $\text{CCSP}(\Gamma)$ is polynomial time solvable if and only if Γ is non-crossing decomposable. Otherwise it is NP-complete.

The equivalence of the two characterizations will be proved in Section 3.

Following [12], we also study the counting problem $\#\text{CCSP}(\Gamma)$ corresponding to $\text{CCSP}(\Gamma)$, in which the objective is to find the number of solutions of a CSP instance that satisfy a global cardinality constraint specified. Creignou et al. [12] proved that if Γ is a constraint language on a 2-element set, say, $\{0, 1\}$, then $\#\text{CCSP}(\Gamma)$ are solvable in polynomial time exactly when $\text{CCSP}(\Gamma)$ is, that is, if every relation from Γ is width-2-affine. Otherwise it is $\text{P}^{\#\text{P}}$ -complete.

We prove that in the general case as well, $\#\text{CCSP}(\Gamma)$ is polynomial time solvable if and only if $\text{CCSP}(\Gamma)$ is. However, in this paper we do not prove a complexity dichotomy, as we do not determine the exact complexity of the hard counting problems. All such problems are NP-hard as Theorems 1.2 and 1.3 show; and we do not claim that the NP-hard cases are actually $\text{P}^{\#\text{P}}$ -hard.

Theorem 1.4. For a constraint language Γ , the problem $\#\text{CCSP}(\Gamma)$ is polynomial time solvable if and only if Γ has a majority polymorphism and a conservative Mal'tsev polymorphism; or, equivalently, if and only if Γ is non-crossing decomposable. Otherwise it is NP-hard.

We also consider the so called *meta-problem* for $\text{CCSP}(\Gamma)$ and $\#\text{CCSP}(\Gamma)$: Suppose set D is fixed. Given a finite constraint language Γ on D , decide whether or not $\text{CCSP}(\Gamma)$ (and $\#\text{CCSP}(\Gamma)$) is solvable in polynomial time. By Theorems 1.2 and 1.4 it suffices to check if Γ has a majority and a conservative Mal'tsev polymorphism. Since the set D is fixed, this can be done by checking, for each possible ternary function with the required properties, whether or not it is a polymorphism of Γ . To check if a ternary operation f is a polymorphism of Γ one just needs for each relation $R \in \Gamma$ to apply f to every triple of tuples in R . This can be done in a time cubic in the total size of relations in Γ .

Theorem 1.5. Let D be a finite set. The meta-problem for $\text{CCSP}(\Gamma)$ and $\#\text{CCSP}(\Gamma)$ is polynomial time solvable.

Note that all the results use the assumption that the set D is fixed (although the complexity of algorithms does not depend on a particular constraint language). Without this assumption the algorithms given in the paper become exponential time, and Theorem 1.5 does not answer if the meta problem is polynomial time solvable if the set D is not fixed, and is a part of the input. The algorithm sketched above is then super-exponential.

2. PRELIMINARIES

Relations and constraint languages. Our notation concerning tuples and relations is fairly standard. The set of all tuples of elements from a set D is denoted by D^n . We denote

tuples in boldface, e.g., \mathbf{a} , and their components by $\mathbf{a}[1], \mathbf{a}[2], \dots$. An n -ary *relation* on set D is any subset of D^n . Sometimes we use instead of relation R the corresponding predicate $R(x_1, \dots, x_n)$. A set of relations on D is called a *constraint language* over D .

For a subset $I = \{i_1, \dots, i_k\} \subseteq \{1, \dots, n\}$ with $i_1 < \dots < i_k$ and an n -tuple \mathbf{a} , by $\text{pr}_I \mathbf{a}$ we denote the *projection of \mathbf{a} onto I* , the k -tuple $(\mathbf{a}[i_1], \dots, \mathbf{a}[i_k])$. The *projection* $\text{pr}_I R$ of R is the k -ary relation $\{\text{pr}_I \mathbf{a} \mid \mathbf{a} \in R\}$. Sometimes we need to emphasize that the unary projections $\text{pr}_1 R, \text{pr}_2 R$ of a binary relation R are sets A and B . We denote this by $R \subseteq A \times B$.

Pairs from equivalence relations play a special role, so such pairs will be denoted by, e.g., $\langle a, b \rangle$. If α is an equivalence relation on a set D then D/α denotes the set of α -classes, and a^α for $a \in D$ denotes the α -class containing a . We say that the equivalence relation α on a set D is *trivial* if D is the only α class.

Constraint Satisfaction Problem with cardinality constraints. Let D be a finite set (throughout the paper we assume it fixed) and Γ a constraint language over D . An instance of the *Constraint Satisfaction Problem* (CSP for short) $\text{CSP}(\Gamma)$ is a pair $\mathcal{P} = (V, \mathcal{C})$, where V is a finite set of *variables* and \mathcal{C} is a set of *constraints*. Every constraint is a pair $C = \langle \mathbf{s}, R \rangle$ consisting of an n_C -tuple \mathbf{s} of variables, called the *constraint scope* and an n_C -ary relation $R \in \Gamma$, called the *constraint relation*. A solution of \mathcal{P} is a mapping $\varphi : V \rightarrow D$ such that for every constraint $C = \langle \mathbf{s}, R \rangle$ the tuple $\varphi(\mathbf{s})$ belongs to R .

A *global cardinality constraint* for a CSP instance \mathcal{P} is a mapping $\pi : D \rightarrow \mathbb{N}$ with $\sum_{a \in D} \pi(a) = |V|$. A solution φ of \mathcal{P} satisfies the cardinality constraint π if the number of variables mapped to each $a \in D$ equals $\pi(a)$. The variant of $\text{CSP}(\Gamma)$ allowing global cardinality constraints will be denoted by $\text{CCSP}(\Gamma)$; the question is, given an instance \mathcal{P} and a cardinality constraint π , whether there is a solution of \mathcal{P} satisfying π .

Example 2.1. If Γ is a constraint language on the 2-element set $\{0, 1\}$ then to specify a global cardinality constraint it suffices to specify the number of ones we want to have in a solution. This problem is also known as the k -ONES(Γ) problem [12].

Example 2.2. Let $\Gamma_{3\text{-COL}}$ be the constraint language on $D = \{0, 1, 2\}$ containing only the binary disequality relation \neq . It is straightforward that $\text{CSP}(\Gamma_{3\text{-COL}})$ is equivalent to the GRAPH 3-COLORABILITY problem. Therefore $\text{CCSP}(\Gamma_{3\text{-COL}})$ is equivalent to the GRAPH 3-COLORABILITY problem in which the question is whether there is a coloring with a prescribed number of vertices colored each color.

Sometimes it is convenient to use arithmetic operations on cardinality constraints. Let $\pi, \pi' : D \rightarrow \mathbb{N}$ be cardinality constraints on a set D , and $c \in \mathbb{N}$. Then $\pi + \pi'$ and $c\pi$ denote cardinality constraints given by $(\pi + \pi')(a) = \pi(a) + \pi'(a)$ and $(c\pi)(a) = c \cdot \pi(a)$, respectively, for any $a \in D$. Furthermore, we extend addition to sets Π, Π' of cardinality vectors in a convolution sense: $\Pi + \Pi'$ is defined as $\{\pi + \pi' \mid \pi \in \Pi, \pi' \in \Pi'\}$.

It is possible to consider an even more general CSP with global cardinality constraints, in which every instance of $\text{CSP}(\Gamma)$ is accompanied with a set of global cardinality constraints, and the question is whether or not there exists a solution of the CSP instance that satisfies one of the cardinality constraints. Sometimes such a set of cardinality constraints can be represented concisely, for example, all constraints π with $\pi(a) = k$. We denote such *extended CSP with global cardinality constraints* corresponding to a constraint language Γ by $\text{ECCSP}(\Gamma)$.

Example 2.3. The problem $\text{ECCSP}(\Gamma_{3\text{-COL}})$ admits a wide variety of questions, e.g. whether a given graph admits a 3-coloring with 25 vertices colored 0, and odd number of vertices colored 1.

However, in our setting (as $|D|$ is a fixed constant and we are investigating polynomial-time solvability) the extended problems are not very interesting from the complexity point of view.

Lemma 2.4. For any constraint language Γ the problem $\text{ECCSP}(\Gamma)$ is Turing reducible to $\text{CCSP}(\Gamma)$.

Proof. Since D is fixed, for any instance \mathcal{P} of $\text{CSP}(\Gamma)$ there are only polynomially many global cardinality constraints. Thus we can try each of the cardinality constraints given in an instance of $\text{ECCSP}(\Gamma)$ in turn. \square

Note that the algorithm in Section 4 for CCSP actually finds the set of all feasible cardinality constraints. Thus ECCSP can be solved in a more direct way than the reduction in Lemma 2.4.

Primitive positive definitions and polymorphisms. Let Γ be a constraint language on a set D . A relation R is *primitive positive (pp-) definable* in Γ if it can be expressed using (a) relations from Γ , (b) conjunction, (c) existential quantifiers, and (d) the binary equality relations (see, e.g. [13]). The set of all relations pp-definable in Γ will be denoted by $\langle\langle\Gamma\rangle\rangle$.

Example 2.5. An important example of pp-definitions that will be used throughout the paper is the *product* of binary relations. Let R, Q be binary relations. Then $R \circ Q$ is the binary relation given by

$$(R \circ Q)(x, y) = \exists z R(x, z) \wedge Q(z, y).$$

In this paper we will need a slightly weaker notion of definability. We say that R is *pp-definable* in Γ *without equalities* if it can be expressed using only items (a)–(c) from above. The set of all relations pp-definable in Γ without equalities will be denoted by $\langle\langle\Gamma\rangle\rangle'$. Clearly, $\langle\langle\Gamma\rangle\rangle' \subseteq \langle\langle\Gamma\rangle\rangle$. The two sets are different only on relations with redundancies. Let R be a (say, n -ary) relation. A *redundancy* of R is a pair i, j of its coordinate positions such that, for any $\mathbf{a} \in R$, $\mathbf{a}[i] = \mathbf{a}[j]$.

Example 2.6. In some cases if a relation R has redundancies, the equality relation is pp-definable in $\{R\}$ without equalities. Let R be a ternary relation on $D = \{0, 1, 2\}$ given by (triples, members of the relation, are written vertically)

$$R = \begin{pmatrix} 0 & 0 & 1 & 1 & 2 & 2 \\ 0 & 0 & 1 & 1 & 2 & 2 \\ 1 & 2 & 0 & 2 & 0 & 1 \end{pmatrix}.$$

Then the equality relation is expressed by $\exists z R(x, y, z)$.

In other cases the equality relation cannot be expressed that easily, but its restriction onto a subset of D can. Let Q be a 4-ary on $D = \{0, 1, 2\}$ given by

$$R = \begin{pmatrix} 0 & 0 & 0 & 2 & 2 & 2 \\ 0 & 0 & 0 & 2 & 2 & 2 \\ 1 & 2 & 0 & 2 & 0 & 1 \\ 0 & 0 & 1 & 1 & 2 & 2 \end{pmatrix}.$$

Then the formula $\exists z, tR(x, y, z, t)$ defines the equality relation on $\{0, 2\}$.

Lemma 2.7. For every constraint language Γ , every $R \in \langle\langle\Gamma\rangle\rangle$ without redundancies belongs to $\langle\langle\Gamma\rangle\rangle'$.

Proof. Consider a pp-definition of R in Γ . Suppose that the definition contains an equality relation on the variables x and y . If none of x and y is bound by an existential quantifier, then the relation R has two coordinates that are always equal, i.e., R is redundant. Thus one of the variables, say x , is bound by an existential quantifier. In this case, replacing x with y everywhere in the definition defines the same relation R and decreases the number of equalities used. Repeating this step, we can arrive to an equality-free definition of R . \square

A *polymorphism* of a (say, n -ary) relation R on D is a mapping $f : D^k \rightarrow D$ for some k such that for any tuples $\mathbf{a}_1, \dots, \mathbf{a}_k \in R$ the tuple

$$f(\mathbf{a}_1, \dots, \mathbf{a}_k) = (f(\mathbf{a}_1[1], \dots, \mathbf{a}_k[1]), \dots, f(\mathbf{a}_1[n], \dots, \mathbf{a}_k[n]))$$

belongs to R . Operation f is a polymorphism of a constraint language Γ if it is a polymorphism of every relation from Γ . There is a tight connection, a *Galois correspondence*, between polymorphisms of a constraint language and relations pp-definable in the language, see [15, 3]. This connection has been extensively exploited to study the ordinary constraint satisfaction problems [21, 9]. Here we do not need the full power of this Galois correspondence, we only need the following folklore result:

Lemma 2.8. If operation f is a polymorphism of a constraint language Γ , then it is also a polymorphism of any relation from $\langle\langle\Gamma\rangle\rangle$, and therefore of any relation from $\langle\langle\Gamma\rangle\rangle'$.

For a (say, n -ary) relation R over a set D and a subset $D' \subseteq D$, by $R|_{D'}$ we denote the relation $\{(a_1, \dots, a_n) \mid (a_1, \dots, a_n) \in R \text{ and } a_1, \dots, a_n \in D'\}$. For a constraint language Γ over D we use $\Gamma|_{D'}$ to denote the constraint language $\{R|_{D'} \mid R \in \Gamma\}$.

If f is a polymorphism of a constraint language Γ over D and $D' \subset D$, then f is not necessarily a polymorphism of $\Gamma|_{D'}$. However, it remains a polymorphism in the following special case. A k -ary polymorphism f is *conservative*, if $f(a_1, \dots, a_k) \in \{a_1, \dots, a_k\}$ for every $a_1, \dots, a_k \in D$. It is easy to see that if f is a conservative polymorphism of Γ , then f is a (conservative) polymorphism of $\Gamma|_{D'}$ for every $D' \subseteq D$.

Polymorphisms help to express many useful properties of relations. A (say, n -ary) relation R is said to be *2-decomposable* if $\mathbf{a} \in R$ if and only if, for any $i, j \in \{1, \dots, n\}$, $\text{pr}_{i,j}\mathbf{a} \in \text{pr}_{i,j}R$, see [1, 20]. Decomposability sometimes is a consequence of the existence of certain polymorphisms. A ternary operation m on a set D is said to be a *majority operation* if it satisfies equations $m(x, x, y) = m(x, y, x) = m(y, x, x) = x$ for all $x, y \in D$. By [1] if a majority operation m is a polymorphism of a constraint language Γ then Γ is 2-decomposable. The converse is not true: there are 2-decomposable relations not having a majority polymorphism. Furthermore, 2-decomposability is not preserved by pp-definitions, thus we cannot expect to characterize it by polymorphisms.

Example 2.9. Consider the disequality relation \neq over the set $D = \{1, 2, 3\}$. Relation \neq is trivially 2-decomposable, since it is binary. Let $R(x, y, z) = \exists q((x \neq q) \wedge (y \neq q) \wedge (z \neq q))$. The binary projections of R are $D \times D$, but R is not $D \times D \times D$: it does not allow that x, y, z are all different.

A binary relation R is said to be *rectangular* if for any $(a, c), (a, d), (b, d) \in R$, the pair (b, c) also belongs to R . Rectangular relations and their generalizations play a very

important role in the study of CSP [7, 8, 19]. A useful way to think about binary rectangular relations is to represent them as *thick mappings*. A binary relation $R \subseteq A \times B$ is called a thick mapping if there are equivalence relations α and β on A and B , respectively, and a one-to-one mapping $\varphi: A/\alpha \rightarrow B/\beta$ (thus, in particular, $|A/\alpha| = |B/\beta|$) such that $(a, b) \in R$ if and only if $b^\beta = \varphi(a^\alpha)$. In this case we shall also say that R is a *thick mapping with respect to α , β , and φ* . Given a thick mapping R the corresponding equivalence relations will be denoted by α_R^1 and α_R^2 . Observe that $\alpha_R^1 = R \circ R^{-1}$ and $\alpha_R^2 = R^{-1} \circ R$; therefore $\alpha_R^1, \alpha_R^2 \in \langle\langle\{R\}\rangle\rangle$. Thick mapping R is said to be *trivial* if both α_R^1 and α_R^2 are the total equivalence relations $(\text{pr}_1 R)^2$ and $(\text{pr}_2 R)^2$. In a graph-theoretical point of view, a thick mapping defines a bipartite graph where every connected component is a complete bipartite graph.

As with decomposability, rectangularity follows from the existence of a certain polymorphism. A ternary operation h is said to be *Mal'tsev* if $h(x, x, y) = h(y, x, x) = y$ for all $x, y \in D$. The first part of following lemma is straightforward, while the second part is folklore

Lemma 2.10. (1) Binary relation R is a thick mapping if and only if it is rectangular.
 (2) If a binary relation R has a Mal'tsev polymorphism then it is rectangular.

A ternary operation h satisfying equations $h(x, x, y) = h(x, y, x) = h(y, x, x) = y$ for all $x, y \in D$ is said to be a *minority operation*. Observe that every Mal'tsev operation is minority, but not the other way round.

Consistency. Let us fix a constraint language Γ on a set D and let $\mathcal{P} = (V, \mathcal{C})$ be an instance of $\text{CSP}(\Gamma)$. A *partial solution* of \mathcal{P} on a set of variables $W \subseteq V$ is a mapping $\psi: W \rightarrow D$ that satisfies every constraint $\langle W \cap \mathbf{s}, \text{pr}_{W \cap \mathbf{s}} R \rangle$ where $\langle \mathbf{s}, R \rangle \in \mathcal{C}$. Here $W \cap \mathbf{s}$ denotes the subtuple of \mathbf{s} consisting of those entries of \mathbf{s} that belong to W , and we consider the coordinate positions of R indexed by variables from \mathbf{s} . Instance \mathcal{P} is said to be *k-consistent* if for any k -element set $W \subseteq V$ and any $v \in V \setminus W$ any partial solution on W can be extended to a partial solution on $W \cup \{v\}$, see [20]. As we only need $k = 2$, all further definitions are given under this assumption.

Any instance $\mathcal{P} = (V, \mathcal{C})$ can be transformed to a 2-consistent instance by means of the standard 2-CONSISTENCY algorithm. This algorithm works as follows. First, for each pair $v, w \in V$ it creates a constraint $\langle (v, w), R_{v,w} \rangle$ where $R_{v,w}$ is the binary relation consisting of all partial solutions ψ on $\{v, w\}$, i.e. $R_{v,w}$ includes pairs $(\psi(v), \psi(w))$. These new constraints are added to \mathcal{C} , let the resulting instance be denoted by $\mathcal{P}' = (V, \mathcal{C}')$. Second, for each pair $v, w \in V$, every partial solution $\psi \in R_{v,w}$, and every $u \in V \setminus \{v, w\}$, the algorithm checks if ψ can be extended to a partial solution of \mathcal{P}' on $\{v, w, u\}$. If not, it updates \mathcal{P}' by removing ψ from $R_{v,w}$. The algorithm repeats this step until no more changes happen.

Lemma 2.11. Let $\mathcal{P} = (V, \mathcal{C})$ be an instance of $\text{CSP}(\Gamma)$.

- (a) The problem obtained from \mathcal{P} by applying 2-CONSISTENCY is 2-consistent;
- (b) On every step of 2-CONSISTENCY for any pair $v, w \in V$ the relation $R_{v,w}$ belongs to $\langle\langle\Gamma\rangle\rangle'$.

Proof. (a) follows from [11].

(b) Since after the first phase of the algorithm every relation $R_{v,w}$ is an intersection of unary and binary projections of relations from Γ , it belongs to $\langle\langle\Gamma\rangle\rangle'$. Then when considering

a pair $v, w \in V$ and $u \in V \setminus \{v, w\}$, the relation $R_{v,w}$ is replaced with $R_{v,w} \cap \text{pr}_{v,w}Q$, where Q is the set of all solution of the current instance on $\{v, w, u\}$. As every relation of the current instance belongs to $\langle\langle\Gamma\rangle\rangle'$, the relation Q is pp-definable in Γ without equalities. Thus the updated relation $R_{v,w}$ also belongs to $\langle\langle\Gamma\rangle\rangle'$. \square

Note that Theorem 2.11(b) implies that any polymorphism of Γ is also a polymorphism of every $R_{v,w}$.

If Γ has a majority polymorphism, by Theorem 3.5 of [20], every 2-consistent problem instance is *globally consistent*, that is every partial solution can be extended to a global solution of the problem. In particular, if \mathcal{P} is 2-consistent, then for any $v, w \in V$, any pair $(a, b) \in R_{v,w}$ can be extended to a solution of \mathcal{P} . The same is true for any $a \in \text{pr}_1 R_{v,w}$ and any $b \in \text{pr}_2 R_{v,w}$.

3. EQUIVALENCE OF THE CHARACTERIZATIONS

In this section, we prove that the two characterizations in Theorems 1.2 and 1.3 are equivalent. Recall that two equivalence relations α and β over the same domain *cross*, if there is an α -class A and a β -class B such that none of $A \setminus B$, $A \cap B$, and $B \setminus A$ is empty. A constraint language Γ is said to be *non-crossing decomposable* if every relation from $\langle\langle\Gamma\rangle\rangle$ is 2-decomposable, every binary relation from $\langle\langle\Gamma\rangle\rangle$ is a thick mapping, and no pair of equivalence relations from $\langle\langle\Gamma\rangle\rangle$ cross.

One of the directions is easy to see.

Lemma 3.1. If Γ has a majority polymorphism m and a conservative Mal'tsev polymorphism h , then Γ is non-crossing decomposable.

Proof. As m is a polymorphism of every relation in $\langle\langle\Gamma\rangle\rangle$, by [1] every relation in $\langle\langle\Gamma\rangle\rangle$ is 2-decomposable. Similarly, every binary relation in $\langle\langle\Gamma\rangle\rangle$ is invariant under the Mal'tsev polymorphism, which, by Lemma 2.10, implies that the binary relation is a thick mapping.

Finally, suppose that there are two equivalence relations $\alpha, \beta \in \langle\langle\Gamma\rangle\rangle$ over the same domain D' that cross. This means that for some $a, b, c \in D'$ we have that $\langle a, b \rangle \in \alpha$, $\langle b, c \rangle \in \beta$, but $\langle a, b \rangle \notin \beta$, $\langle b, c \rangle \notin \alpha$. Let h be a Mal'tsev polymorphism of Γ , and consider $d = h(a, b, c)$. First of all, as h is conservative, $d \in \{a, b, c\}$. Then, since h preserves α and β ,

$$h \begin{pmatrix} a & b & c \\ a & a & c \end{pmatrix} = \begin{pmatrix} d \\ c \end{pmatrix} \in \alpha, \quad h \begin{pmatrix} a & b & c \\ a & c & c \end{pmatrix} = \begin{pmatrix} d \\ a \end{pmatrix} \in \beta,$$

which is impossible. \square

To prove the other direction of the equivalence, we need to construct the two polymorphisms. The following definition will be useful for this purpose.

Definition 3.2. Given a constraint language Γ , we say that $(a|bc)$ is true if $\langle\langle\Gamma\rangle\rangle$ contains an equivalence relation α with

$$\langle a, a \rangle, \langle b, b \rangle, \langle c, c \rangle, \langle b, c \rangle, \langle c, b \rangle \in \alpha \text{ and } \langle a, b \rangle, \langle a, c \rangle, \langle b, a \rangle, \langle c, a \rangle \notin \alpha.$$

In other words, the domain of the equivalence relation contains all three elements, b and c are in the same class, but a is in a different class.

Lemma 3.3. Let Γ be a non-crossing decomposable constraint language over D . For every $a, b, c \in D$, at most one of $(a|bc)$, $(b|ac)$, $(c|ab)$ is true.

Proof. Suppose that, say, both $(a|bc)$ and $(b|ac)$ are true; let α, β be the corresponding equivalence relations from $\langle\langle\Gamma\rangle\rangle$. Let D_α, D_β be the domains of α and β , respectively. We can consider D_α and D_β as unary relations, and they are in $\langle\langle\Gamma\rangle\rangle$. Therefore, $D' = D_\alpha \cap D_\beta$ and $\alpha' = \alpha \cap (D' \times D')$ and $\beta' = \beta \cap (D' \times D')$ are also in $\langle\langle\Gamma\rangle\rangle$. As is easily seen, α' and β' are over the same domain D' and they cross, a contradiction. \square

Lemma 3.4. Let Γ be a non-crossing decomposable constraint language over D . Let $R \in \langle\langle\Gamma\rangle\rangle$ be a binary relation such that $(a, a'), (b, b'), (c, c') \in R$, but $(p, q) \notin R$ for some $p \in \{a, b, c\}$ and $q \in \{a', b', c'\}$. Then

$$\begin{aligned} (a|bc) &\iff (a'|b'c') \\ (b|ac) &\iff (b'|a'c') \\ (c|ab) &\iff (c'|a'b'). \end{aligned}$$

Proof. If, say, a and b are in the same α_R^1 class, then c has to be in a different class. Since R is a thick mapping, this means that a' and b' are in the same α_R^2 class and c' is in a different class. Therefore, both $(c|ab)$ and $(c'|a'b')$ are true, and by Lemma 3.3, none of the other statements can be true. Therefore, we can assume that a, b, c are in different α_R^1 classes, and hence a', b', c' are in different α_R^2 classes.

Suppose that $(a|bc)$ is true; let $\alpha \in \langle\langle\Gamma\rangle\rangle$ be the corresponding equivalence relation. Consider the relation $R' = \alpha \circ R$, that is, $R'(x, y) = \exists z(\alpha(x, z) \wedge R(z, y))$, which has to be a thick mapping. Let $\beta_{R'}^1, \beta_{R'}^2$ be the equivalence relations of R' . Observe that $\beta_{R'}^1, \beta_{R'}^2 \in \langle\langle R' \rangle\rangle \subseteq \langle\langle\Gamma\rangle\rangle$. We claim that $\langle b', c' \rangle \in \beta_{R'}^2$ and $\langle a', b' \rangle \notin \beta_{R'}^2$, showing that $(a'|b'c')$ is true. It is clear that $\langle b', c' \rangle \in \beta_{R'}^2$: as $(b, b), (b, c) \in \alpha$, we have $(b, b'), (b, c') \in R'$. To get a contradiction suppose that $\langle a', b' \rangle \in \beta_{R'}^2$. Since $(a, a'), (b, b') \in R'$ and R' is a thick mapping, the pairs $(a, b'), (b, a')$ has to belong to R' as well. That is, there are x, y such that $\langle a, x \rangle, \langle b, y \rangle \in \alpha$ and $(x, b'), (y, a') \in R$. Now equivalence relation α shows that $(y|ax)$ is true (since $\langle a, x \rangle, \langle b, y \rangle \in \alpha$ and $\langle b, a \rangle \notin \alpha$) and equivalence relation α_R^1 shows that $(x|ay)$ is true (since $(a, a'), (y, a') \in R$ shows $\langle a, y \rangle \in \alpha_R^1$, $(b, b'), (x, b') \in R$ shows $\langle b, x \rangle \in \alpha_R^1$, and we know that $\langle a, b \rangle \notin \alpha_R^1$). By Lemma 3.3, $(y|ax)$ and $(x|ay)$ cannot be both true, a contradiction. \square

Let Γ be a non-crossing decomposable language over D . Let $\text{minor}(a, b, c)$ be a if a, b, c are all different or all the same, otherwise let it be the value that appears only once among a, b, c . Similarly, let $\text{major}(a, b, c)$ be a if a, b, c are all different or all the same, otherwise let it be the value that appears more than once among a, b, c . Because of Lemma 3.3, the following two functions are well defined:

$$\begin{aligned} m(a, b, c) &= \begin{cases} b & \text{if } (a|bc) \text{ is true,} \\ a & \text{if } (b|ac) \text{ is true,} \\ b & \text{if } (c|ab) \text{ is true,} \\ \text{major}(a, b, c) & \text{if none of } (a|bc), (b|ac), (c|ab) \text{ is true,} \end{cases} \\ h(a, b, c) &= \begin{cases} a & \text{if } (a|bc) \text{ is true,} \\ b & \text{if } (b|ac) \text{ is true,} \\ c & \text{if } (c|ab) \text{ is true,} \\ \text{minor}(a, b, c) & \text{if none of } (a|bc), (b|ac), (c|ab) \text{ is true,} \end{cases} \end{aligned}$$

Lemma 3.5. Operations m and h are conservative majority and minority operations, respectively, and $\langle\langle\Gamma\rangle\rangle$ is invariant under m and h .

Proof. It is clear that m and h are conservative. To show that m is a majority operation, by definition of major, it is sufficient to consider the case when one of $(a|bc), (b|ac), (c|ab)$ is true. If $b = c$ (resp., $a = c, a = b$), then only $(a|bc)$ (resp., $(b|ac), (c|ab)$) can be true, which means that $m(a, b, c) = b$ (resp., a, b), as required. A similar argument shows that h is a minority function.

Since every relation in $\langle\langle\Gamma\rangle\rangle$ is 2-decomposable, it is sufficient to show that every binary relation in $\langle\langle\Gamma\rangle\rangle$ is invariant under m and h . We show invariance under h , the proof is similar for m . Let $R \in \langle\langle\Gamma\rangle\rangle$ be a binary relation, which is a thick mapping by assumption. Take $(a, a'), (b, b'), (c, c') \in R$. If a, b, c are in the same α_R^1 -class then a', b', c' are in the same α_R^2 -class. Since $h(a, b, c) \in \{a, b, c\}$ and $h(a', b', c') \in \{a', b', c'\}$, it follows that $(h(a, b, c), h(a', b', c')) \in R$. If a, b, c are not all in the same α_R^1 -class, and one of $(a|bc), (b|ac), (c|ab)$ is true, then by Lemma 3.4, the corresponding statement from $(a'|b'c'), (b'|a'c'), (c'|a'b')$ is also true. Now the pair $(h(a, b, c), h(a', b', c'))$ has to be one of $(a, a'), (b, b'), (c, c')$, hence it is in R . If none of $(a|bc), (b|ac), (c|ab)$ is true, then a, b, c are in different α_R^1 -classes. Moreover, if none of $(a'|b'c'), (b'|a'c'), (c'|a'b')$ is true, then a', b', c' are in different α_R^2 -classes. Therefore $(h(a, b, c), h(a', b', c')) = (\text{minor}(a, b, c), \text{minor}(a', b', c')) = (a, a') \in R$. \square

Remark 3.6. Interestingly, Lemma 3.5 (along with Lemma 3.1) gives more than just the existence of a majority polymorphism, and a conservative Mal'tsev polymorphism. The operations m and h are both conservative, and h is a minority operation, not just a Mal'tsev one. Therefore, we have that a constraint language has a majority and conservative Mal'tsev polymorphisms if and only if it has a majority and minority polymorphisms, both conservative.

The following two consequences of having a conservative Mal'tsev polymorphism will be used in the algorithm.

Lemma 3.7. Let Γ be a constraint language having a conservative Mal'tsev polymorphism. Let $R, R' \in \langle\langle\Gamma\rangle\rangle$ be two nontrivial thick mappings such that $\text{pr}_2 R = \text{pr}_1 R'$. Then $R \circ R'$ is also non-trivial.

Proof. Since R, R' are nontrivial, there are values a, a', b, b' such that $\langle a, a' \rangle \notin \alpha_R^1, \langle b, b' \rangle \notin \alpha_{R'}^2$. If $R \circ R'$ is trivial, then $(a, b'), (a, b), (a', b) \in R \circ R'$, which means that there are (not necessarily distinct) values x, y, z such that

$$\begin{aligned} (a, x) &\in R, (x, b') \in R', \\ (a, y) &\in R, (y, b) \in R', \\ (a', z) &\in R, (z, b) \in R'. \end{aligned}$$

Let m be a conservative Mal'tsev polymorphism. Let $q = m(x, y, z) \in \{x, y, z\}$. Applying m on the pairs above, we get that $(a', q) \in R$ and $(q, b') \in R'$. It is not possible that $q \in \{x, y\}$, since this would mean that $\langle a, a' \rangle \in \alpha_R^1$. It is not possible that $q = z$ either, since that would imply $\langle b, b' \rangle \in \alpha_{R'}^2$, a contradiction. \square

Recall that if α and β are equivalence relations on the same set S , then $\alpha \vee \beta$ denotes the smallest (in terms of the number of pairs it contains) equivalence relation on S containing both α and β . It is well known that if α and β have a Mal'tsev polymorphism then they *permute*, that is, $\alpha \circ \beta = \beta \circ \alpha = \alpha \vee \beta$.

Lemma 3.8. Let Γ be a finite constraint language having a conservative Mal'tsev polymorphism. Let $\alpha, \beta \in \langle\langle\Gamma\rangle\rangle$ be two nontrivial equivalence relations on the same set S . Then $\alpha \vee \beta$ is also non-trivial.

Proof. Since $\alpha \vee \beta = \alpha \circ \beta$, by Lemma 3.7, it is non-trivial. \square

We will also need the following observation.

Lemma 3.9. Let Γ be a finite constraint language having a majority polymorphism and a conservative Mal'tsev polymorphism. Let $\alpha, \beta \in \langle\langle\Gamma\rangle\rangle$ be two equivalence relations on the same set S . Then $\alpha \vee \beta = \alpha \cup \beta$.

Proof. By Lemma 3.1 α and β are non-crossing. Then if $\langle a, b \rangle \in \alpha$ and $\langle b, c \rangle \in \beta$ then $\langle b, c \rangle, \langle a, c \rangle \in \alpha$. The result follows. \square

4. ALGORITHM

In this section we fix a constraint language Γ that has conservative majority and minority polymorphisms. We present a polynomial-time algorithm for solving $\text{CCSP}(\Gamma)$ and $\#\text{CCSP}(\Gamma)$ in this case.

4.1. Prerequisites. In this subsection we prove several properties of instances of $\text{CCSP}(\Gamma)$ and $\#\text{CCSP}(\Gamma)$ that will be very instrumental for our algorithms. First of all, we show that every such instance can be supposed binary, that is, that every its constraint is imposed only on two variables. Then we introduce a graph corresponding to such an instance, and show that if this graph is disconnected then a solution to the whole problem can be obtained by combining arbitrarily solutions for the connected components. Finally, if the graph is connected, the set of possible values for each variable can be subdivided into several subsets, so that if the variable takes a value from one of the subsets, then each of the remaining variables is forced to take values from a particular subset of the corresponding partition.

Observe that if a constraint language Γ satisfies the conditions of Theorem 1.2 then by Remark 3.6 the constraint language Γ' obtained from Γ by adding all unary relations also satisfies the conditions of Theorem 1.2. Indeed, Γ has conservative majority and minority polymorphisms that are also polymorphisms of Γ' . Therefore we will assume that Γ contains all unary relations. It will also be convenient to assume that Γ contains all the binary relations from $\langle\langle\Gamma\rangle\rangle$.

Let Γ be a constraint language and let $\mathcal{P} = (V, \mathcal{C})$ be a 2-consistent instance of $\text{CSP}(\Gamma)$. By $\text{bin}(\mathcal{P})$ we denote the instance (V, \mathcal{C}') such that \mathcal{C}' is the set of all constraints of the form $\langle\langle v, w \rangle, R_{v,w} \rangle$ where $v, w \in V$ and $R_{v,w}$ is the set of all partial solutions on $\{v, w\}$.

Lemma 4.1. Let Γ be a constraint language with a majority polymorphism. Then if \mathcal{P} is a 2-consistent instance of $\text{CSP}(\Gamma)$ then $\text{bin}(\mathcal{P})$ has the same solutions as \mathcal{P} .

Proof. Let us denote by R, R' the $|V|$ -ary relations consisting of all solutions of \mathcal{P} and $\text{bin}(\mathcal{P})$, respectively. Relations R, R' are pp-definable in Γ without equalities, and $R \subseteq R'$. To show that $R = R'$ we use the result from [20] stating that, since Γ has a majority polymorphism for any $v, w \in V$ and any $(a, b) \in R_{v,w}$ we have $(a, b) \in \text{pr}_{v,w}R$, i.e. $\text{pr}_{v,w}R = \text{pr}_{v,w}R'$. Since R is 2-decomposable, if $\mathbf{a} \in R'$, that is, $\text{pr}_{v,w}\mathbf{a} \in R_{v,w} = \text{pr}_{v,w}R$ for all $v, w \in V$, then $\mathbf{a} \in R$. \square

Let $\mathcal{P} = (V, \mathcal{C})$ be an instance of $\text{CSP}(\Gamma)$. Applying algorithm 2-CONSISTENCY we may assume that \mathcal{P} is 2-consistent. By the assumption about Γ , all constraint relations from \mathcal{P} are 2-decomposable, and $\text{bin}(\mathcal{P})$ has the same solutions as \mathcal{P} itself. Therefore, replacing \mathcal{P} with $\text{bin}(\mathcal{P})$, if necessary, every constraint of \mathcal{P} can be assumed to be binary.

Let constraints of \mathcal{P} be $\langle (v, w), R_{vw} \rangle$ for each pair of different $v, w \in V$. Let $\mathcal{S}_v, v \in V$, denote the set of $a \in D$ such that there is a solution φ of \mathcal{P} such that $\varphi(v) = a$. By [20], \mathcal{P} is globally consistent, therefore, $\mathcal{S}_v = \text{pr}_1 R_{vw}$ for any $w \in V, w \neq v$.

Constraint $C_{vw} = \langle (v, w), R_{vw} \rangle$ is said to be *trivial* if $R_{vw} = \mathcal{S}_v \times \mathcal{S}_w$, otherwise it is said to be *non-trivial*. The *graph of \mathcal{P}* , denoted $G(\mathcal{P})$, is a graph with vertex set V and edge set $E = \{vw \mid v, w \in V \text{ and } C_{vw} \text{ is non-trivial}\}$.

The 2-consistency of \mathcal{P} implies, in particular, the following simple property.

Lemma 4.2. By the 2-consistency of \mathcal{P} , for any $u, v, w \in V$, $R_{uv} \subseteq R_{uw} \circ R_{vw}$.

Therefore, by Lemma 3.7, the graph $G(\mathcal{P})$ is transitive, i.e., every connected component is a clique.

If $G(\mathcal{P})$ is not connected, every combination of solutions for its connected components give rise to a solution of the entire problem. More precisely, let V_1, \dots, V_k be the connected components of $G(\mathcal{P})$, and let $\mathcal{P}_{|V_i}$ denote the instance (V_i, \mathcal{C}_i) where \mathcal{C}_i includes all the constraints $\langle (v, w), R_{vw} \rangle$ for which $v, w \in V_i$. We will use the following observation.

Lemma 4.3. Let $\varphi_1, \dots, \varphi_k$ be solutions of $\mathcal{P}_{|V_1}, \dots, \mathcal{P}_{|V_k}$. Then the mapping $\varphi : V \rightarrow D$ such that $\varphi(v) = \varphi_i(v)$ whenever $v \in V_i$ is a solution of \mathcal{P} .

Proof. We need to check that all constraints of \mathcal{P} are satisfied. Consider $C_{vw} = \langle (v, w), R_{vw} \rangle$. If $v, w \in V_i$ for a certain i , then $(\varphi(v), \varphi(w)) = (\varphi_i(v), \varphi_i(w)) \in R_{vw}$ since φ_i is a solution to \mathcal{P}_i . If v, w belong to different connected components, then R_{vw} is trivial, and so C_{vw} is satisfied. \square

Suppose that $G(\mathcal{P})$ is connected and fix $v \in V$. In this case, the graph is a clique, and therefore for any $w \in V$ the constraint C_{vw} is non-trivial. Note that due to 2-consistency, every relation $\alpha_{R_{vw}}^1$ for $w \in V \setminus \{v\}$ is over the same set \mathcal{S}_v . Set $\eta_v = \bigvee_{w \in V \setminus \{v\}} \alpha_{R_{vw}}^1$; as every $\alpha_{R_{vw}}$ is non-trivial, Lemma 3.8 implies that η_v is non-trivial.

Lemma 4.4. Suppose $G(\mathcal{P})$ is connected. Equivalence relations η_v and $\alpha_{R_{vw}}^1$ (for any $w \in V \setminus \{v\}$) are non-trivial.

Lemma 4.5. Suppose $G(\mathcal{P})$ is connected.

- (1) For any $v, w \in V$ there is a one-to-one correspondence ψ_{vw} between \mathcal{S}_v / η_v and \mathcal{S}_w / η_w such that for any solution φ of \mathcal{P} if $\varphi(v) \in A \in \mathcal{S}_v / \eta_v$, then $\varphi(w) \in \psi_{vw}(A) \in \mathcal{S}_w / \eta_w$.
- (2) The mappings ψ_{vw} are consistent, i.e. for any $u, v, w \in V$ we have $\psi_{uw}(x) = \psi_{vw}(\psi_{uv}(x))$ for every x .

Proof. (1) Let R_{vw} be a thick mapping with respect to a mapping ϱ , and $\alpha = \alpha_{R_{vw}}^1$, $\alpha' = \alpha_{R_{vw}}^2$. Recall that ϱ is a one-to-one mapping from \mathcal{S}_v/α to \mathcal{S}_w/α' . Suppose, to obtain a contradiction, that ϱ does not induce a one-to-one mapping between \mathcal{S}_v/η_v and \mathcal{S}_w/η_w . Then without loss of generality there are $a, b \in \mathcal{S}_v$ such that $\langle a, b \rangle \in \eta_v$, but for certain $a', b' \in \mathcal{S}_w$ we have $(a, a'), (b, b') \in R_{vw}$ and $\langle a', b' \rangle \notin \eta_w$. Since $\alpha' \subseteq \eta_w$, $\langle a', b' \rangle \notin \alpha'$, hence $\langle a, b \rangle \notin \alpha$. By Lemma 3.9 there is $u \in V$ such that R_{vu} is a thick mapping with respect to some β, β' and $\langle a, b \rangle \in \beta$. Therefore for some $c \in \mathcal{S}_u$ we have $(a, c), (b, c) \in R_{vu}$. Since $R_{vu} \subseteq R_{vw} \circ R_{wu}$, there exist $d_1, d_2 \in \mathcal{S}_w$ satisfying the conditions $(a, d_1), (b, d_2) \in R_{vw}$ and $(d_1, c), (d_2, c) \in R_{wu}$. The first pair of inclusions imply that $\langle a', d_1 \rangle, \langle b', d_2 \rangle \in \alpha'$, while the second one implies that $\langle d_1, d_2 \rangle \in \eta_w$. Since $\alpha' \subseteq \eta_w$, we obtain $\langle a', b' \rangle \in \eta_w$, a contradiction.

(2) If for some $u, v, w \in V$ there is a class $A \in \mathcal{S}_u/\eta_u$ such that $\psi_{vw}(\psi_{uv}(A)) \neq \psi_{uw}(A)$ then $R_{uw} \not\subseteq R_{uv} \circ R_{vw}$, a contradiction. \square

Fix a variable v_0 of \mathcal{P} and take a η_{v_0} -class A . Let $\mathcal{P}_A = (V, \mathcal{C}_A)$ denote the problem instance over the same variables, where for every $v, w \in V$ the set \mathcal{C}_A includes the constraint $\langle (v, w), R_{vw}^A \rangle$ with $R_{vw}^A = R_{vw} \cap (\psi_{v_0v}(A) \times \psi_{v_0w}(A))$.

Lemma 4.6. Problem \mathcal{P}_A belongs to $\text{CCSP}(\Gamma)$.

Proof. It suffices to show that $R_{vw}^A \in \Gamma$ for any $v, w \in V$. By Lemma 2.11 $R_{vw} \in \langle \langle \Gamma \rangle \rangle$, and as we assumed that Γ contains all binary relations from $\langle \langle \Gamma \rangle \rangle$, we have $R_{vw} \in \Gamma$. By the assumption made, all unary relations including $\psi_{v_0v}(A)$ and $\psi_{v_0w}(A)$ belong to Γ . Therefore relation R_{vw}^A is pp-definable in Γ , and, as a binary relation, belongs to it. \square

4.2. Algorithm: The decision problem. We split the algorithm into two parts. Algorithm CARDINALITY (Figure 1) just ensures 2-consistency and initializes a recursive process. The main part of the work is done by EXT-CARDINALITY (Figure 2).

Algorithm EXT-CARDINALITY solves the more general problem of computing the set of all cardinality constraints π that can be satisfied by a solution of \mathcal{P} . Thus it can be used to solve directly CSP with extended global cardinality constraints, see Preliminaries.

The algorithm considers three cases. Step 2 handles the trivial case when the instance consists of a single variable and there is only one possible value it can be assigned. Otherwise, we decompose the instance either by partitioning the variables or by partitioning the domain of the variables. If $G(\mathcal{P})$ is not connected, then the satisfying assignments of \mathcal{P} can be obtained from the satisfying assignments of the connected components. Thus a cardinality constraint π can be satisfied if it arises as the sum $\pi_1 + \dots + \pi_k$ of cardinality constraints such that the i -th component has a solution satisfying π_i . Instead of considering all such sums (which would not be possible in polynomial time), we follow the standard dynamic programming approach of going through the components one by one, and determining all possible cardinality constraints that can be satisfied by a solution for the first i components (Step 3).

If the graph $G(\mathcal{P})$ is connected, then we fix a variable v_0 and go through each class A of the partition η_{v_0} (Step 4). If v_0 is restricted to A , then this implies a restriction for every other variable w . We recursively solve the problem for the restricted instance \mathcal{P}_A arising for each class A ; if constraint π can be satisfied, then it can be satisfied for one of the restricted instances.

The correctness of the algorithm follows from the discussion above. The only point that has to be verified is that the instance remains 2-consistent after the recursion. This is obvious if we recurse on the connected components (Step 3). In Step 4, 2-consistency follows from the fact that if $(a, b) \in R_{vw}$ can be extended by $c \in S_u$, then in every subproblem either these three values satisfy the instance restricted to $\{v, w, u\}$ or a, b, c do not appear in the domain of v, w, u , respectively.

To show that the algorithm runs in polynomial time, observe first that every step of the algorithm (except the recursive calls) can be done in polynomial time. Here we use that D is fixed, hence the size of the set Π is polynomially bounded. Thus we only need to bound the size of the recursion tree. If we recurse in Step 3, then we produce instances whose graphs are connected, thus it cannot be followed by recursing again in Step 3. In Step 4, the domain of every variable is decreased: by Lemma 4.4, η_w is nontrivial for any variable w . Thus in any branch of the recursion tree, recursion in Step 4 can occur at most $|D|$ times, hence the depth of the recursion tree is $O(|D|)$. As the number of branches is polynomial in each step, the size of the recursion tree is polynomial.

INPUT: An instance $\mathcal{P} = (V, \mathcal{C})$ of CCSP(Γ) with a cardinality constraint π
 OUTPUT: YES if \mathcal{P} has a solution satisfying π , NO otherwise

Step 1. **apply** 2-CONSISTENCY to \mathcal{P}
Step 2. **set** $\Pi := \text{EXT-CARDINALITY}(\mathcal{P})$
Step 3. **if** $\pi \in \Pi$ **output** YES
 else output NO

Figure 1: Algorithm CARDINALITY.

4.3. Solving the counting problem. In this section we observe that algorithm CARDINALITY can be modified so that it also solves counting CSPs with global cardinality constraints, provided Γ satisfies the conditions of Theorem 1.2.

The counting algorithm works very similar to algorithm CARDINALITY, except that instead of determining the set of satisfiable cardinality constraints, it keeps track of the number of solutions that satisfy every cardinality constraint possible. It considers the same 3 cases. In the trivial case of a problem with one variable and one possible value for this variable, the algorithm assigns 1 to the cardinality constraint satisfied by the only solution of the problem and 0 to all other cardinality constraints. In the case of disconnected graph $G(\mathcal{P})$ if a cardinality constraint can be represented in the form $\pi = \pi_1 + \dots + \pi_k$, then solutions on the connected components of $G(\mathcal{P})$ satisfying π_1, \dots, π_k , respectively, contribute the product of their numbers into the number of solutions satisfied by π . We again use the dynamic programming approach, and, for each i compute the number of solutions on $V_1 \cup \dots \cup V_i$ satisfying every possible cardinality constraint. Observe, that the set of cardinality constraints considered is also changed dynamically, as the number of variables grows. Finally, if $G(\mathcal{P})$ is connected, then the different restrictions have disjoint sets of solutions, hence the numbers of solutions are computed independently.

5. DEFINABLE RELATIONS, CONSTANT RELATIONS, AND THE COMPLEXITY OF CCSP

We present two reductions that will be crucial for the proofs in Section 6. In Section 5.1, we show that adding relations that are pp-definable (without equalities) does not make the problem harder, while in Section 5.2, we show the same for unary constant relations.

INPUT: A 2-consistent instance $\mathcal{P} = (V, \mathcal{C})$ of $\text{CCSP}(\Gamma)$
 OUTPUT: The set of cardinality constraints π such that \mathcal{P} has a solution that satisfies π

```

Step 1.  construct the graph  $G(\mathcal{P}) = (V, E)$ 
Step 2.  if  $|V| = 1$  and the domain of this variable is a singleton  $\{a\}$  then do
Step 2.1    set  $\Pi := \{\pi\}$  where  $\pi(x) = 0$  except for  $\pi(a) = 1$ 
Step 3.  else if  $G(\mathcal{P})$  is disconnected and  $G_1 = (V_1, E_1), \dots, G_k = (V_k, E_k)$  are
its connected components then do
Step 3.1    set  $\Pi := \{\pi\}$  where  $\pi : D \rightarrow \mathbb{N}$  is given by  $\pi(a) = 0$  for  $a \in D$ 
Step 3.2    for  $i = 1$  to  $k$  do
Step 3.2.1    set  $\Pi := \Pi + \text{EXT-CARDINALITY}(\mathcal{P}|_{V_i})$ 
endfor
endif
Step 4.  else do
Step 4.1    for each  $v \in V$  find  $\eta_v$ 
Step 4.2    fix  $v_0 \in V$  and set  $\Pi := \emptyset$ 
Step 4.3    for each  $\eta_{v_0}$ -class  $A$  do
Step 4.3.1    set  $\mathcal{P}_A := (V, \mathcal{C}_A)$  where for every  $v, w \in V$  the set  $\mathcal{C}_A$  includes
the constraint  $\langle (v, w), R_{vw} \cap (\psi_{v_0 v}(A) \times \psi_{v_0 w}(A)) \rangle$ 
Step 4.3.2    set  $\Pi := \Pi \cup \text{EXT-CARDINALITY}(\mathcal{P}_A)$ 
endfor
enddo
Step 4.  output  $\Pi$ 

```

Figure 2: Algorithm EXT-CARDINALITY.

INPUT: An instance $\mathcal{P} = (V, \mathcal{C})$ of $\#\text{CCSP}(\Gamma)$ with a cardinality constraint π
 OUTPUT: The number of solutions of \mathcal{P} that satisfy π

```

Step 1.  apply 2-CONSISTENCY to  $\mathcal{P}$ 
Step 2.  set  $\varrho := \#\text{EXT-CARDINALITY}(\mathcal{P})$ 
        %  $\varrho(\pi')$  is the number of solutions of  $\mathcal{P}$  satisfying cardinality constraint  $\pi'$ 
Step 3.  output  $\varrho(\pi)$ 

```

Figure 3: Algorithm #CARDINALITY.

5.1. Definable relations and the complexity of cardinality constraints.

Theorem 5.1. Let Γ be a constraint language and R a relation pp-definable in Γ without equalities. Then $\text{CCSP}(\Gamma \cup \{R\})$ is polynomial time reducible to $\text{CCSP}(\Gamma)$.

Proof. We proceed by induction on the structure of pp-formulas. The base case of induction is given by $R \in \Gamma$. We need to consider two cases.

CASE 1. $R(x_1, \dots, x_n) = R_1(x_1, \dots, x_n) \wedge R_2(x_1, \dots, x_n)$.

Observe that by introducing ‘fictitious’ variables for predicates R_1, R_2 we may assume that both relations involved have the same arity. A reduction from $\text{CCSP}(\Gamma \cup \{R\})$ to $\text{CCSP}(\Gamma)$ is trivial: in a given instance of the first problem replace each constraint of the form $\langle (v_1, \dots, v_n), R \rangle$ with two constraints $\langle (v_1, \dots, v_n), R_1 \rangle$ and $\langle (v_1, \dots, v_n), R_2 \rangle$.

CASE 2. $R(x_1, \dots, x_n) = \exists x R'(x_1, \dots, x_n, x)$.

INPUT: A 2-consistent instance $\mathcal{P} = (V, \mathcal{C})$ of $\#CCSP(\Gamma)$
 OUTPUT: Function ϱ that assigns to every cardinality constraint π with $\sum_{a \in D} \pi(a) = |V|$, the number $\varrho(\pi)$ of solutions of \mathcal{P} that satisfy π

Step 1. **construct** the graph $G(\mathcal{P}) = (V, E)$
Step 2. **if** $|V| = 1$ and the domain of this variable is a singleton $\{a\}$ **then do**
Step 2.1 **set** $\varrho(\pi) := 1$ where $\pi(x) = 0$ except for $\pi(a) = 1$, and $\varrho(\pi') := 0$ for all $\pi' \neq \pi$
 with $\sum_{x \in D} \pi'(x) = 1$
Step 3. **else if** $G(\mathcal{P})$ is disconnected and $G_1 = (V_1, E_1), \dots, G_k = (V_k, E_k)$ are its
 connected components **then do**
Step 3.1 **set** $\Pi := \{\pi\}$ where $\pi : D \rightarrow \mathbb{N}$ is given by $\pi(a) = 0$ for $a \in D$, $\varrho(\pi) := 1$ for $\pi \in \Pi$
Step 3.2 **for** $i = 1$ **to** k **do**
Step 3.2.1 **set** $\Pi' := \{\pi : D \rightarrow \mathbb{N} \mid \sum_{a \in D} \pi(a) = |V_i|\}$ and $\varrho' := \#EXT-CARDINALITY(\mathcal{P}|_{V_i})$
Step 3.2.2 **set** $\Pi'' := \{\pi : D \rightarrow \mathbb{N} \mid \sum_{a \in D} \pi(a) = |V_1| + \dots + |V_i|\}$, $\varrho''(\pi) := 0$ for $\pi \in \Pi''$
Step 3.2.3 **for each** $\pi \in \Pi$ **and** $\pi' \in \Pi'$ **set** $\varrho''(\pi + \pi') := \varrho''(\pi + \pi') + \varrho(\pi) \cdot \varrho'(\pi')$
Step 3.2.4 **set** $\Pi := \Pi''$, $\varrho := \varrho''$
 endfor
 endif
Step 4. **else do**
Step 4.1 **for each** $v \in V$ **find** η_v
Step 4.2 **fix** $v_0 \in V$ **and set** $\varrho(\pi) := 0$ for π with $\sum_{a \in D} \pi(a) = |V|$
Step 4.3 **for each** η_{v_0} -class A **do**
Step 4.3.1 **set** $\mathcal{P}_A := (V, \mathcal{C}_A)$ where for every $v, w \in V$ the set \mathcal{C}_A includes the constraint
 $\langle (v, w), R_{vw} \cap (\psi_{v_0 v}(A) \times \psi_{v_0 w}(A)) \rangle$
Step 4.3.2 **set** $\varrho' := \#EXT-CARDINALITY(\mathcal{P}_A)$
Step 4.3.3 **set** $\varrho(\pi) := \varrho(\pi) + \varrho'(\pi)$
 endfor
 enddo
Step 4. **output** ϱ

Figure 4: Algorithm $\#EXT-CARDINALITY$.

Let $\mathcal{P} = (V, \mathcal{C})$ be a $CCSP(\Gamma \cup \{R\})$ instance. Without loss of generality let C_1, \dots, C_q be the constraints that involve R . Instance \mathcal{P}' of $CCSP(\Gamma)$ is constructed as follows.

- Variables: Replace every variable v from V with a set W_v of variables of size $q|D|$ and introduce a set of $|D|$ variables for each constraint involving R . More formally,

$$W = \bigcup_{v \in V} W_v \cup \{w_1, \dots, w_q\} \cup \bigcup_{i=1}^q \{w_i^1, \dots, w_i^{|D|-1}\}.$$

- Non- R constraints: For every $C_i = \langle (v_1, \dots, v_\ell), Q \rangle$ with $i > q$, introduce all possible constraints of the form $\langle (u_1, \dots, u_\ell), Q \rangle$, where $u_j \in W_{v_j}$ for $j \in \{1, \dots, \ell\}$.
- R constraints: For every $C_i = \langle (v_1, \dots, v_\ell), R \rangle$, $i \leq q$, introduce all possible constraints of the form $\langle (u_1, \dots, u_\ell, w_i), R' \rangle$, where $u_j \in W_{v_j}$, $j \in \{1, \dots, \ell\}$.

CLAIM 1. If \mathcal{P} has a solution satisfying cardinality constraint π then \mathcal{P}' has a solution satisfying the cardinality constraint $\pi' = q|D| \cdot \pi + q$.

Let φ be a solution of \mathcal{P} satisfying π . It is straightforward to verify that the following mapping ψ is a solution of \mathcal{P}' and satisfies π' :

- for each $v \in V$ and each $u \in W_v$ set $\psi(u) = \varphi(v)$;

- for each w_i , where $C_i = \langle (v_1, \dots, v_n), R \rangle$, set $\psi(w_i)$ to be a value such that $(\varphi(v_1), \dots, \varphi(v_n), \psi(w_i)) \in R'$.
- for each $i \leq q$ and $j \leq |D| - 1$ set $\psi(w_i^j)$ to be such that $\{\psi(w_i), \psi(w_i^1), \dots, \psi(w_i^{|D|-1})\} = D$.

CLAIM 2. If \mathcal{P}' has a solution ψ satisfying the cardinality constraint $\pi' = q|D| \cdot \pi + q$, then \mathcal{P} has a solution satisfying constraint π .

Let $a \in D$ and $U_a(\psi) = \psi^{-1}(a) = \{u \in W \mid \psi(u) = a\}$. Observe first that if $\varphi : V \rightarrow D$ is a mapping such that $U_{\varphi(v)}(\psi) \cap W_v \neq \emptyset$ for every $v \in V$ (i.e., $\psi(v') = \varphi(v)$ for at least one variable $v' \in W_v$), then φ satisfies all the constraints of \mathcal{P} . Indeed, consider a constraint $C = \langle \mathbf{s}, Q \rangle$ of \mathcal{P} where $Q \neq R$. Let $\mathbf{s} = (v_1, \dots, v_\ell)$. For every v_i , there is a $v'_i \in W_{v_i}$ such that $\varphi(v_i) = \psi(v'_i)$. By the way \mathcal{P}' is defined, it contains a constraint $C' = \langle \mathbf{s}', Q \rangle$ where $\mathbf{s}' = (v'_1, \dots, v'_\ell)$. Now the fact that ψ satisfies C' immediately implies that φ satisfies C : $(\varphi(v_1), \dots, \varphi(v_\ell)) = (\psi(v'_1), \dots, \psi(v'_\ell)) \in Q$. The argument is similar if $Q = R$.

We show that it is possible to construct such a φ that also satisfies the cardinality constraint π . Since $|W_v| = q|D|$, for any $a \in D$ with $\pi(a) \neq 0$, even if set $U_a(\psi)$ contains all $q|D|$ variables of the form w_i and w_i^j , it has to intersect at least $\pi(a)$ sets W_v (as $(\pi(a)-1)q|D| + q|D| < \pi'(a) = \pi(a) \cdot q|D| + q$). Consider the bipartite graph $G = (T_1 \cup T_2, E)$, where T_1, T_2 is a bipartition and

- T_1 is the set of variables V ;
 - T_2 is constructed from the set D of values by taking $\pi(a)$ copies of each value $a \in D$;
 - edge (v, a') , where a' is a copy of a from T_2 , belongs to E if and only if $W_v \cap U_a(\psi) \neq \emptyset$.
- Note that $|T_1| = |T_2|$ and a perfect matching $E' \subseteq E$ corresponds to a required mapping φ : $\varphi(v) = a$ if $(v, a') \in E'$ for some copy a' or a .

Take any subset $S \subseteq T_2$, let S contains some copies of a_1, \dots, a_s . Then by the observation above, S has at least $\pi(a_1) + \dots + \pi(a_s)$ neighbours in T_1 . Since S contains at most $\pi(a_i)$ copies of a_i ,

$$\pi(a_1) + \dots + \pi(a_s) \geq |S|.$$

By Hall's Theorem on perfect matchings in bipartite graphs, G has a perfect matching, concluding the proof that the required φ exists. \square

5.2. Constant relations and the complexity of cardinality constraints. Let D be a set, and let $a \in D$. The *constant relation* C_a is the unary relation that contains only one tuple, (a) . If a constraint language Γ over D contains all the constant relations, then they can be used in the corresponding constraint satisfaction problem to force certain variables to take some fixed values. The goal of this section is to show that for any constraint language Γ the problem $\text{CCSP}(\Gamma \cup \{C_a \mid a \in D\})$ is polynomial time reducible to $\text{CCSP}(\Gamma)$. For the ordinary decision CSP such a reduction exists when Γ does not have unary polymorphisms that are not permutations, see [9].

We make use of the notion of multi-valued morphisms, a generalization of homomorphisms, that in a different context has appeared in the literature for a while (see, e.g. [26]) under the guise hyperoperation. Let R be a (say, n -ary) relation on a set D , and let f be a mapping from D to 2^D , the powerset of D . Mapping f is said to be a *multi-valued morphism* of R if for any tuple $(a_1, \dots, a_n) \in R$ the set $f(a_1) \times \dots \times f(a_n)$ is a subset of R . Mapping f is a multi-valued morphism of a constraint language Γ if it is a multi-valued morphism of every relation in Γ . For a multi-valued morphism f and set $A \subseteq D$, we define

$f(A) := \bigcup_{a \in A} f(a)$. The product of two multi-valued morphisms f_1 and f_2 is defined by $(f_1 \circ f_2)(a) := f_1(f_2(a))$ for every $a \in D$. We denote by f^i the i -th power of f , with the convention that f^0 maps a to $\{a\}$ for every $a \in A$.

Theorem 5.2. Let Γ be a finite constraint language over a set D . Then $\text{CCSP}(\Gamma \cup \{C_a \mid a \in D\})$ is polynomial time reducible to $\text{CCSP}(\Gamma)$.

Proof. Let $D = \{d_1, \dots, d_k\}$ and $a = d_1$. We show that $\text{CCSP}(\Gamma \cup \{C_a\})$ is polynomial time reducible to $\text{CCSP}(\Gamma)$. This clearly implies the result. We make use of the following multi-valued morphism gadget $\text{MVM}(\Gamma, n)$ (i.e. a CSP instance). Observe that it is somewhat similar to the *indicator problem* [22].

- The set of variables is $V(n) = \bigcup_{i=1}^k V_{d_i}$, where V_{d_i} contains $n^{|D|+1-i}$ elements. All sets V_{d_i} are assumed to be disjoint.
- The set of constraints is constructed as follows: For every (say, r -ary) $R \in \Gamma$ and every $(a_1, \dots, a_r) \in R$ we include all possible constraints of the form $\langle (v_1, \dots, v_r), R \rangle$ where $v_i \in V_{a_i}$ for $i \in \{1, \dots, r\}$.

Now, given an instance $\mathcal{P} = (V, \mathcal{C})$ of $\text{CCSP}(\Gamma \cup \{C_a\})$, we construct an instance $\mathcal{P}' = (V', \mathcal{C}')$ of $\text{CCSP}(\Gamma)$.

- Let $W \subseteq V$ be the set of variables v , on which the constant relation C_a is imposed, that is, \mathcal{C} contains the constraint $\langle (v), C_a \rangle$. Set $n = |V|$. The set V' of variables of \mathcal{P}' is the disjoint union of the set $V(n)$ of variables of $\text{MVM}(\Gamma, n)$ and $V \setminus W$.
- The set \mathcal{C}' of constraints of \mathcal{P}' consists of three parts:
 - (a) \mathcal{C}'_1 , the constraints of $\text{MVM}(\Gamma, n)$;
 - (b) \mathcal{C}'_2 , the constraints of \mathcal{P} that do not include variables from W ;
 - (c) \mathcal{C}'_3 , for any constraint $\langle (v_1, \dots, v_n), R \rangle \in \mathcal{C}$ whose scope contains variables constrained by C_a (without loss of generality let v_1, \dots, v_ℓ be such variables), \mathcal{C}'_3 contains all constraints of the form $\langle (w_1, \dots, w_\ell, v_{\ell+1}, \dots, v_n), R \rangle$, where $w_1, \dots, w_\ell \in V_a$.

We show that \mathcal{P} has a solution satisfying a cardinality constraint π if and only if \mathcal{P}' has a solution satisfying cardinality constraint π' given by

$$\pi'(d_i) = \begin{cases} \pi(a) + (|V_a| - |W|), & \text{if } i = 1, \\ \pi(d_i) + |V_{d_i}|, & \text{otherwise.} \end{cases}$$

Suppose that \mathcal{P} has a right solution φ . Then a required solution for \mathcal{P}' is given by

$$\psi(v) = \begin{cases} \varphi(v), & \text{if } v \in V \setminus W, \\ d_i, & \text{if } v \in V_{d_i}. \end{cases}$$

It is straightforward that ψ is a solution to \mathcal{P}' and that it satisfies π' .

Suppose that \mathcal{P}' has a solution ψ that satisfies π' . Since $\pi'(a) = \pi(a) + n^{|D|} - |W| \geq n^{|D|} - n > |V' \setminus V_a|$, there is $v \in V_a$ such that $\psi(v) = a$. Thus the assignment

$$\varphi(v) = \begin{cases} \psi(v), & \text{if } v \in V \setminus W, \\ a & \text{if } v \in W \end{cases}$$

is a satisfying assignment \mathcal{P} , but it might not satisfy π . Our goal is to show that \mathcal{P}' has a solution ψ , where φ obtained this way satisfies π . Observe that what we need is that ψ assigns value d_i to exactly $\pi'(d_i) - |V_{d_i}|$ variables of $V \setminus W$.

CLAIM 1. Mapping f taking every $d_i \in D$ to the set $\{\psi(v) \mid v \in V_{d_i}\}$ is a multi-valued morphism of Γ .

Indeed, let $(a_1, \dots, a_n) \in R$, R is an $(n\text{-ary})$ relation from Γ . Then by the construction of $\text{MVM}(\Gamma, n)$ the instance contains all the constraints of the form $\langle (v_1, \dots, v_n), R \rangle$ with $v_i \in V_{a_i}$, $i \in \{1, \dots, n\}$. Therefore,

$$\begin{aligned} & \{\psi(v_1) \mid v_1 \in V_{a_1}\} \times \dots \times \{\psi(v_n) \mid v_n \in V_{a_n}\} \\ &= f(a_1) \times \dots \times f(a_n) \subseteq R. \end{aligned}$$

CLAIM 2. Let f be the mapping defined in Claim 1. Then f^* defined by $f^*(b) := f(b) \cup \{b\}$ for every $b \in D$ is also a multi-valued morphism of Γ .

We show that for every $d_i \in D$, there is an $n_i \geq 1$ such that $d_i \in f^j(d_i)$ for every $j \geq n_i$. Taking the maximum $n = \max_{1 \leq i \leq k} n_i$ of all these integers, we get that $d_i \in f^{n+1}(d_i)$ and $f(d_i) \subseteq f^{n+1}(d_i)$ (since $d_i \in f^n(d_i)$) for every i , proving the claim.

The proof is by induction on i . If $d_i \in f(d_i)$, then we are done as we can set $n_i = 1$ (note that this is always the case for $i = 1$, since we observed above that ψ has to assign value d_1 to a variable of V_{d_1}). So let us suppose that $d_i \notin f(d_i)$. Let $D_i = \{d_1, \dots, d_i\}$ and let $g_i : D_i \rightarrow 2^{D_i}$ be defined by $g_i(d_j) := f(d_j) \cap D_i$. Observe that $g_i(d_j)$ is well-defined, i.e., $g_i(d_j) \neq \emptyset$: the set V_{d_j} contains $n^{|D|+1-j} \geq n^{|D|+1-i}$ variables, while the number of variables which are assigned by ψ values outside D_i is $\sum_{\ell=i+1}^k \pi'(d_\ell) \leq n + \sum_{\ell=i+1}^k n^{|D|+1-\ell} < n^{|D|+1-i}$.

Let $T := \bigcup_{\ell \geq 1} g_i^\ell(d_i)$. We claim that $d_i \in T$. Suppose that $d_i \notin T$. By the definition of T and the assumption $d_i \notin f(d_i)$, for every $b \in T \cup \{d_i\}$, the variables in V_b can have values only from T and from $D \setminus D_i$. The total number of variables in V_b , $b \in T \cup \{d_i\}$ is $\sum_{b \in T \cup \{d_i\}} n^{|D|+1-b}$, while the total cardinality constraint of the values from $T \cup (D \setminus D_i)$ is

$$\begin{aligned} \sum_{b \in T \cup (D \setminus D_i)} \pi'(b) &= n + \sum_{b \in T} n^{|D|+1-b} + \sum_{\ell=i+1}^k n^{|D|+1-\ell} \\ &< \sum_{b \in T} n^{|D|+1-b} + n^{|D|+1-i} = \sum_{b \in T \cup \{d_i\}} n^{|D|+1-b}, \end{aligned}$$

a contradiction. Thus $d_i \in T$, that is, there is a value $j < i$ such that $d_j \in f(d_i)$ and $d_i \in f^s(d_j)$ for some $s \geq 1$. By the induction hypothesis, $d_j \in f^n(d_j)$ for every $n \geq n_j$, thus we have that $d_i \in f^n(d_i)$ if n is at least $n_i := 1 + n_j + s$. This concludes the proof of Claim 2.

Let D^+ (resp., D^-) be the set of those values $d_i \in D$ that ψ assigns to more than (resp., less than) $\pi'(i) - |V_{d_i}|$ variables of $V \setminus W$. It is clear that if $|D^+| = |D^-| = 0$, then φ obtained from ψ satisfies π . Furthermore, if $|D^+| = 0$, then $|D^-| = 0$ as well. Thus suppose that $D^+ \neq \emptyset$ and let $S := \bigcup_{b \in D^+, s \geq 1} f^s(b)$.

CLAIM 3. $S \cap D^- \neq \emptyset$.

Suppose $S \cap D^- = \emptyset$. Every $b \in S \subseteq D \setminus D^-$ is assigned by ψ to at least $\pi'(b) - |V_b|$ variables in $V \setminus W$, implying that ψ assigns every such b to at most $|V_b|$ variables in the gadget $\text{MVM}(\Gamma, n)$. Thus the total number of variables in the gadget having value from S is at most $\sum_{b \in S} |V_b|$; in fact, it is strictly less than that since D^+ is not empty. By the definition of S , only values from S can be assigned to variables in V_b for every $b \in S$. However, the total number of these variables is exactly $\sum_{b \in S} |V_b|$, a contradiction.

By Claim 3, there is a value $d^- \in S \cap D^-$, which means that there is a $d^+ \in D^+$ and a sequence of distinct values $b_0 = d^+$, $b_1, \dots, b_\ell = d^-$ such that $b_{i+1} \in f(b_i)$ for every $0 \leq i < \ell$. Let $v \in V \setminus W$ be an arbitrary variable having value d^+ . We modify ψ the following way:

- (1) The value of v is changed from d^+ to d^- .
- (2) For every $0 \leq i < \ell$, one variable in V_{b_i} with value b_{i+1} is changed to b_i . Note that since $b_{i+1} \in f(b_i)$ and $b_{i+1} \neq b_i$ such a variable exists.

Note that these changes do not modify the cardinalities of the values: the second step increases only the cardinality of $b_0 = d^+$ (by one) and decreases only the cardinality of $b_\ell = d^-$ (by one). We have to argue that the transformed assignment still satisfies the constraints of \mathcal{P}' . Since $d^- \in f^\ell(d^+)$, the multi-valued morphism f^* of Claim 2 implies that changing d^+ to d^- on a single variable and not changing anything else also gives a satisfying assignment. To see that the second step does not violate the constraints, observe first that constraints of type (b) are not affected and constraints of type (c) cannot be violated (since variables in V_{d_1} are changed only to d_1 , and there is already at least one variable with value d_1 in V_{d_1}). To show that constraints of type (a) are not affected, it is sufficient to show that the mapping f' described by the gadget after the transformation is still a multi-valued morphism. This can be easily seen as $f'(b) \subseteq f(b) \cup \{b_i\} = f^*(b)$, where f^* is the multi-valued morphism of Claim 2.

Thus the modified assignment is still a solution of \mathcal{P}' satisfying π' . It is not difficult to show that repeating this operation, in a finite number of steps we reach a solution where $D^+ = D^- = \emptyset$, i.e., every value $b \in D(b)$ appears exactly $\pi'(b) - |V(b)|$ times on the variables of $V \setminus W$. As observed above, this implies that \mathcal{P} has a solution satisfying π . \square

6. HARDNESS

Now we prove that if constraint language Γ does not satisfy the conditions of Theorem 1.2 then CCSP(Γ) is NP-complete. First, we can easily simulate the restriction to a subset of the domain by setting to 0 the cardinality constraint on the unwanted values:

Lemma 6.1. For any constraint language Γ over a set D and any $D' \subseteq D$, the problem CCSP($\Gamma|_{D'}$) is polynomial time reducible to CCSP(Γ).

Proof. For an instance $\mathcal{P}' = (V, \mathcal{C}')$ of CCSP($\Gamma|_{D'}$) with a global cardinality constraint $\pi' : D' \rightarrow \mathbb{N}$ we construct an instance $\mathcal{P} = (V, \mathcal{C})$ of CCSP(Γ) such that for each $\langle \mathbf{s}, R|_{D'} \rangle \in \mathcal{C}'$ we include $\langle \mathbf{s}, R \rangle$ into \mathcal{C} . The cardinality constraint π' is replaced with $\pi : D \rightarrow \mathbb{N}$ such that $\pi(a) = \pi'(a)$ for $a \in D'$, and $\pi(a) = 0$ for $a \in D \setminus D'$. It is straightforward that \mathcal{P} has a solution satisfying π if and only if \mathcal{P}' has a solution satisfying π' . \square

Suppose now that a constraint language Γ does not satisfy the conditions of Theorem 1.2. Then one of the following cases takes place: (a) $\langle \langle \Gamma \rangle \rangle$ contains a binary relation that is not a thick mapping; or (b) $\langle \langle \Gamma \rangle \rangle$ contains a crossing pair of equivalence relations; or (c) $\langle \langle \Gamma \rangle \rangle$ contains a relation which is not 2-decomposable. Since relations occurring in cases (a), (b) are not redundant, and relations that may occur in case (c) can be assumed to be not redundant, by Lemma 2.7 $\langle \langle \Gamma \rangle \rangle$ can be replaced with $\langle \langle \Gamma \rangle \rangle'$. Furthermore, by Theorem 5.2 all the constant relations can be assumed to belong to Γ . We consider these three cases in turn. Furthermore, by a minimality argument, we can assume that if Γ is over D , then for every $S \subset D$, constraint language $\Gamma|_S$ satisfies the requirements of Theorem 1.2.

One of the NP-complete problems we will reduce to $\text{CCSP}(R)$ is the BIPARTITE INDEPENDENT SET problem (or BIS for short). In this problem given a connected bipartite graph G with bipartition V_1, V_2 and numbers k_1, k_2 , the goal is to verify if there exists an independent set S of G such that $|S \cap V_1| = k_1$ and $|S \cap V_2| = k_2$. The NP-completeness of the problem follows for example from [10], which shows the NP-completeness of the complement problem under the name CONSTRAINED MINIMUM VERTEX COVER. It is easy to see that the problem is hard even for graphs containing no isolated vertices. By representing the edges of a bipartite graph with the relation $R = \{(a, c), (a, d), (b, d)\}$, we can express the problem of finding a bipartite independent set. Value b (resp., a) represents selected (resp., unselected) vertices in V_1 , while value c (resp., d) represents selected (resp., unselected) vertices in V_2 . With this interpretation, the only combination that relation R excludes is that two selected vertices are adjacent. By Lemma 2.10, if a binary relation is not a thick mapping, then it contains something very similar to R . However, some of the values a, b, c , and d might coincide and the relation might contain further tuples such as (c, d) . Thus we need a delicate case analysis to show that the problem is NP-hard for binary relations that are not thick mappings.

Lemma 6.2. Let R be a binary relation which is not a thick mapping. Then $\text{CCSP}(\{R\})$ is NP-complete.

Proof. Since R is not a thick mapping, there are $(a, c), (a, d), (b, d) \in R$ such that $(b, c) \notin R$. By Lemma 6.1 the problem $\text{CCSP}(R')$, where $R' = R|_{\{a, b, c, d\}}$, is polynomial time reducible to $\text{CCSP}(R)$. Replacing R with R' if necessary we can assume that R is a relation over $D = \{a, b, c, d\}$ (note that some of those elements can be equal). We suppose that R is a ‘smallest’ relation that is not a thick mapping, that is, for any R' definable in $\{R\}$ with $R' \subset R$, the relation R' is a thick mapping, and for any subset D' of D the restriction of R onto D' is a thick mapping.

Let $B = \{x \mid (a, x) \in R\}$. Since $B(x) = \exists y R(y, x) \wedge C_a(y)$, the unary relation B is definable in R . If $B \neq \text{pr}_2 R$, by setting $R'(x, y) = R(x, y) \wedge B(y)$ we get a binary relation R' that is not a thick mapping. Thus by the minimality of R , we may assume that $(a, x) \in R$ for any $x \in \text{pr}_2 R$, and symmetrically, $(y, d) \in R$ for any $y \in \text{pr}_1 R$.

CASE 1. $|\{a, b, c, d\}| = 4$.

We claim that $|\text{pr}_1 R| = |\text{pr}_2 R| = 2$. Suppose, without loss of generality that $x \in \{a, b\}$ appears in $\text{pr}_2 R$. If $(b, x) \in R$, then $(a, x), (b, x), (a, c) \in R$ and $(b, c) \notin R$. Therefore the restriction $R|_{\{a, b, c\}}$ is not a thick mapping, contradicting the minimality of R . Otherwise $(a, x), (a, d), (b, d) \in R$ while $(b, x) \notin R$. Hence $R|_{\{a, b, d\}}$ is not a thick mapping. Thus we have $\text{pr}_1 R = \{a, b\}$ and $\text{pr}_2 R = \{c, d\}$.

Let $G = (V, E), V_1, V_2, k_1, k_2$ be an instance of BIS. Construct an instance $\mathcal{P} = (V, \mathcal{C})$ by including into \mathcal{C} , for every edge (v, w) of G , the constraint $\langle (v, w), R \rangle$, and defining a cardinality constraint as $\pi(a) = |V_1| - k_1, \pi(b) = k_1, \pi(c) = k_2, \pi(d) = |V_2| - k_2$. It is straightforward that for any solution φ of \mathcal{P} the set $S_\varphi = \{v \in V \mid \varphi(v) \in \{b, c\}\}$ is an independent set, $S_\varphi \cap V_1 = \{v \mid \varphi(v) = b\}$, $S_\varphi \cap V_2 = \{v \mid \varphi(v) = c\}$. Set S_φ satisfies the required conditions if and only if φ satisfies π . Conversely, for any independent set $S \subseteq V$ mapping φ given by

$$\varphi_S(v) = \begin{cases} a, & \text{if } v \in V_1 \setminus S, \\ b, & \text{if } v \in V_1 \cap S, \\ c, & \text{if } v \in V_2 \cap S, \\ d, & \text{if } v \in V_2 \setminus S, \end{cases}$$

is a solution of \mathcal{P} that satisfies π if and only if $|S \cap V_1| = k_1$ and $|S \cap V_2| = k_2$.

CASE 2. $|\{a, b, c, d\}| = 2$.

Then R is a binary relation with 3 tuples in it over a 2-element set. By [12] $\text{CCSP}(R)$ is NP-complete.

Thus in the remaining cases, we can assume that $|\{a, b, c, d\}| = 3$, and therefore $\{a, b\} \cap \{c, d\} \neq \emptyset$.

CLAIM 1. One of the projections $\text{pr}_1 R$ or $\text{pr}_2 R$ contains only 2 elements.

Let us assume the converse. Let $\text{pr}_2 R = \{c, d, x\}$, $x \in \{a, b\}$ (as R is over a 3-element set). We consider two cases. Suppose first $c \notin \{a, b\}$ (implying $d \in \{a, b\}$). If $(b, x) \notin R$, then the restriction of R onto $\{a, b\}$ contains $(a, d), (b, d), (a, x)$, but does not contain (b, x) . Thus $R|_{\{a, b\}}$ is not a thick mapping, a contradiction with the minimality assumption. If $(b, x) \in R$ then the set $B = \{a, b\} = \{x \mid (b, x) \in R\}$ is definable in R . Observe that $R'(x, y) = R(x, y) \wedge B(x)$ is not a thick mapping and definable in R . A contradiction with the choice of R .

Now suppose that $d \notin \{a, b\}$ (implying $c \in \{a, b\}$). If $(b, x) \in R$, then the restriction $R|_{\{a, b\}}$ is not a thick mapping, as $(a, c), (a, x), (b, x) \in R$ and $(b, c) \notin R$. Otherwise let $(b, x) \notin R$. By the assumption made $|\text{pr}_1 R| = 3$, that is, $d \in \text{pr}_1 R$. We consider 4 cases depending on whether (d, c) and (d, x) are contained in R . If $(d, c), (d, x) \notin R$, then, as $a \in \{x, c\}$, the relation $R|_{\{a, d\}}$ is not a thick mapping (recall that $(d, d) \in R$). If $(d, c), (d, x) \in R$, then we can restrict R on $\{d, b\}$ (note that $b \in \{c, x\}$). Finally, if $(d, c) \in R, (d, x) \notin R$ [or $(d, x) \in R, (d, c) \notin R$], then the relation $B = \{d, c\}$ [respectively, $B = \{d, x\}$] is definable in R . It remains to observe that $R'(x, y) = R(x, y) \wedge B(x)$ is not a thick mapping. This concludes the proof of the claim.

Thus we can assume that one of the projections $\text{pr}_1 R$ or $\text{pr}_2 R$ contains only 2 elements. Without loss of generality, let $\text{pr}_1 R = \{a, b\}$. In the remaining cases, we assume $\text{pr}_2 R = \{c, d, x\}$, where $x \in \{a, b\}$ and x may not be present.

CASE 3. Either

- $c \notin \{a, b\}$ (SUBCASE 3A), or
- $d \notin \{a, b\}$ and $(b, x) \notin R$ (SUBCASE 3B).

In this case, given an instance $G = (V, E), V_1, V_2, k_1, k_2$ of BIS, we construct an instance $\mathcal{P} = (V', \mathcal{C})$ of $\text{CCSP}(R)$ as follows.

- $V' = V_2 \cup \bigcup_{w \in V_1} V^w$, where all the sets V_2 and V^w , $w \in V_1$, are disjoint, and $|V^w| = 2|V|$.
- For any $(u, w) \in E$ the set \mathcal{C} contains all constraints of the form $\langle (v, w), R \rangle$ where $v \in V^u$.
- The cardinality constraint π is given by the following rules:
 - Subcase 3a:
 - $\pi(c) = k_2$, $\pi(a) = (|V_1| - k_1) \cdot 2|V|$, $\pi(b) = k_1 \cdot 2|V| + (|V_2| - k_2)$ if $d = b$, and
 - $\pi(c) = k_2$, $\pi(a) = (|V_1| - k_1) \cdot 2|V| + (|V_2| - k_2)$, $\pi(b) = k_1 \cdot 2|V|$ if $d = a$.
 - Subcase 3b:
 - $\pi(d) = |V_2| - k_2$, $\pi(a) = (|V_1| - k_1) \cdot 2|V|$, $\pi(b) = k_1 \cdot 2|V| + k_2$ if $c = b$, and
 - $\pi(d) = |V_2| - k_2$, $\pi(a) = (|V_1| - k_1) \cdot 2|V| + k_2$, $\pi(b) = k_1 \cdot 2|V|$ if $c = a$.

If G has a required independent set S , then consider a mapping $\varphi : V' \rightarrow D$ given by

$$\varphi(v) = \begin{cases} a, & \text{if } v \in V^w \text{ and } w \in V_1 \setminus S, \\ b, & \text{if } v \in V^w \text{ and } w \in V_1 \cap S, \\ c, & \text{if } v \in V_2 \cap S, \\ d, & \text{if } v \in V_2 \setminus S, \end{cases}$$

For any $\langle (u, v), R \rangle \in \mathcal{C}$, $u \in V^w$, either $w \notin S$ or $v \notin S$. In the first case $\varphi(u) = a$ and so $(\varphi(u), \varphi(v)) \in R$. In the second case $\varphi(u) = b$ and $\varphi(v) = d$. Again, $(\varphi(u), \varphi(v)) \in R$. Finally it is straightforward that φ satisfies the cardinality constraint π .

Suppose that \mathcal{P} has a solution φ that satisfies π . Since $\text{pr}_1 R = \{a, b\}$ and we can assume that G has no isolated vertices, for any $u \in V^w$, $w \in V_1$, we have $\varphi(u) \in \{a, b\}$. Also if for some $u \in V^w$ it holds that $\varphi(u) = b$ and $\varphi(v) = c$ for $v \in V_2$ then $(w, v) \notin E$. We include into $S \subseteq V$ all vertices $w \in V_1$ such that there is $u \in V^w$ with $\varphi(u) = b$. By the choice of the cardinality of V^w and $\pi(b)$ there are at least k_1 such vertices. In Subcase 3a, we include in S all vertices $v \in V_2$ with $\varphi(v) = c$. There are exactly k_2 vertices like this, and by the observation above S is an independent set. In Subcase 3b, we include in S all vertices $v \in V_2$ with $\varphi(v) \in \{a, b\}$. By the choice of $\pi(d)$, there are at least k_2 such vertices. To verify that S is an independent set it suffices to recall that in this case $(b, x) \notin R$, and so $(b, a), (b, b) \notin R$.

CASE 4. $d \notin \{a, b\}$ and $(b, x) \in R$.

In this case $\{c, x\} = \{a, b\}$ and $(a, c), (a, x), (b, x) \in R$ while $(b, c) \notin R$. Therefore $R|_{\{a, b\}}$ is not a thick mapping. A contradiction with the choice of R . \square

Next we show hardness in the case when there is a crossing pair of equivalence relations. With a simple observation, we can obtain a binary relation that is not a thick mapping and then apply Lemma 6.2.

Lemma 6.3. Let α, β be a crossing pair of equivalence relations. Then $\text{CCSP}(\{\alpha, \beta\})$ is NP-complete.

Proof. Let α, β be equivalence relations on the same domain D . This means that there are $a, b, c \in D$ such that $\langle a, c \rangle \in \alpha \setminus \beta$ and $\langle c, b \rangle \in \beta \setminus \alpha$. Let $\alpha' = \alpha|_{\{a, b, c\}}$ and $\beta' = \beta|_{\{a, b, c\}}$. Clearly,

$$\begin{aligned} \alpha' &= \{(a, a), (b, b), (c, c), (a, c), (c, a)\}, \\ \beta' &= \{(a, a), (b, b), (c, c), (b, c), (c, b)\}. \end{aligned}$$

By Lemma 6.1, $\text{CCSP}(\{\alpha', \beta'\})$ is polynomial-time reducible to $\text{CCSP}(\{\alpha, \beta\})$. Consider the relation $R = \alpha' \circ \beta'$, that is, $R(x, y) = \exists z(\alpha'(x, z) \wedge \beta'(z, y))$. We have that $\text{CCSP}(R)$ is reducible to $\text{CCSP}(\{\alpha', \beta'\})$ and

$$R = \{(a, a), (b, b), (c, c), (a, c), (c, a), (b, c), (c, b), (a, b)\}.$$

Observe that R is not a thick mapping and by Lemma 6.2, $\text{CCSP}(R)$ is NP-complete. \square

Finally, we prove hardness in the case when there is a relation that is not 2-decomposable. An example of such a relation is a ternary Boolean affine relation $x + y + z = c$ for $c = 0$ or $c = 1$. The CSP with global cardinality constraints for this relation is NP-complete by [12]. Our strategy is to obtain such a relation from a relation that is not 2-decomposable. However, as in Lemma 6.2, we have to consider several cases.

We start with the following simple lemma:

Lemma 6.4. Let α be an equivalence relation on a set D and $a \in D$. Then $a^\alpha \in \langle\langle\alpha, C_a\rangle\rangle'$.

Proof. The unary relation $R(x) = \exists y((\langle x, y \rangle \in \alpha) \wedge C_a(y))$ is equal to a^α and is clearly in $\langle\langle\alpha, C_a\rangle\rangle'$. \square

Lemma 6.5. Let R be a relation that is not 2-decomposable. Then $\text{CCSP}(\{R\})$ is NP-complete.

Proof. We can assume that every binary relation in $\langle\langle\{R\} \cup \{C_a \mid a \in D\}\rangle\rangle'$ is a thick mapping, and no pair of equivalence relations from this set cross, otherwise the problem is NP-complete by Theorems 5.1, 5.2, and Lemmas 6.2, 6.3. Furthermore, we can choose R to be the ‘smallest’ non-2-decomposable relation in the sense that every relation obtained from R by restricting on a proper subset of D is 2-decomposable, and every relation $R' \in \langle\langle\{R\} \cup \{C_a \mid a \in D\}\rangle\rangle'$ that either have smaller arity, or $R' \subset R$, is also 2-decomposable.

We claim that relation R is ternary. Indeed, it cannot be binary by assumptions made about it. Suppose that $\mathbf{a} \notin R$ is a tuple such that $\text{pr}_{ij}\mathbf{a} \in \text{pr}_{ij}R$ for any i, j . Let

$$R'(x, y, z) = \exists x_4, \dots, x_n (R(x, y, z, x_4, \dots, x_n) \wedge C_{\mathbf{a}[4]}(x_4) \wedge \dots \wedge C_{\mathbf{a}[n]}(x_n)).$$

It is straightforward that $(\mathbf{a}[1], \mathbf{a}[2], \mathbf{a}[3]) \notin R'$, while, since any proper projection of R is 2-decomposable, $\text{pr}_{\{2, \dots, n\}}\mathbf{a} \in \text{pr}_{\{2, \dots, n\}}R$, $\text{pr}_{\{1, 3, \dots, n\}}\mathbf{a} \in \text{pr}_{\{1, 3, \dots, n\}}R$, $\text{pr}_{\{1, 2, 4, \dots, n\}}\mathbf{a} \in \text{pr}_{\{1, 2, 4, \dots, n\}}R$, implying $(\mathbf{a}[2], \mathbf{a}[3]) \in \text{pr}_{23}R'$, $(\mathbf{a}[1], \mathbf{a}[3]) \in \text{pr}_{13}R'$, $(\mathbf{a}[1], \mathbf{a}[2]) \in \text{pr}_{12}R'$, respectively. Thus R' is not 2-decomposable, a contradiction with assumptions made.

Let $(a, b, c) \notin R$ and $(a, b, d), (a, e, c), (f, b, c) \in R$, and let $B = \{a, b, c, d, e, f\}$. As $R|_B$ is not 2-decomposable, we should have $R = R|_B$.

If $R_{12} = \text{pr}_{12}R$ is a thick mapping with respect to some equivalence relations η_{12}, η_{21} (see p. 2), $R_{13} = \text{pr}_{13}R$ is a thick mapping with respect to η_{13}, η_{31} , and $R_{23} = \text{pr}_{23}R$ is a thick mapping with respect to η_{23}, η_{32} , then $\langle a, f \rangle \in \eta_{12} \cap \eta_{13}$, $\langle b, e \rangle \in \eta_{21} \cap \eta_{23}$, and $\langle c, d \rangle \in \eta_{31} \cap \eta_{32}$. Let the corresponding classes of $\eta_{12} \cap \eta_{13}$, $\eta_{21} \cap \eta_{23}$, and $\eta_{31} \cap \eta_{32}$ be B_1, B_2 , and B_3 , respectively. Then $B_1 = \text{pr}_1R$, $B_2 = \text{pr}_2R$, $B_3 = \text{pr}_3R$. Indeed, if one of these equalities is not true, since by Lemma 6.4 the relations B_1, B_2, B_3 are pp-definable in R without equalities, the relation $R'(x, y, z) = R(x, y, z) \wedge B_1(x) \wedge B_2(y) \wedge B_3(z)$ is pp-definable in R and the constant relations, is smaller than R , and is not 2-decomposable.

Next we show that $(a, g) \in \text{pr}_{12}R$ for all $g \in \text{pr}_2R$. If there is g with $(a, g) \notin \text{pr}_{12}R$ then setting $C(y) = \exists z(\text{pr}_{12}R(z, y) \wedge C_a(z))$ we have $b, e \in C$ and $C \neq \text{pr}_2R$. Thus $R'(x, y, z) = R(x, y, z) \wedge C(y)$ is smaller than R and is not 2-decomposable. The same is true for b and pr_2R , and for c and pr_3R . Since every binary projection of R is a thick mapping this implies that $\text{pr}_{12}R = \text{pr}_1R \times \text{pr}_2R$, $\text{pr}_{23}R = \text{pr}_2R \times \text{pr}_3R$, and $\text{pr}_{13}R = \text{pr}_1R \times \text{pr}_3R$.

For each $i \in \{1, 2, 3\}$ and every $p \in \text{pr}_iR$, the relation $R_i^p(x_j, x_k) = \exists x_i(R(x_1, x_2, x_3) \wedge C_p(x_i))$, where $\{j, k\} = \{1, 2, 3\} \setminus \{i\}$, is definable in R and therefore is a thick mapping with respect to, say, η_{ij}^p, η_{ik}^p . Our next step is to show that R can be chosen such that η_{ij}^p does not depend on the choice of p and i .

If one of these relations, say, R_1^p , equals $\text{pr}_2R \times \text{pr}_3R$, while another one, say R_1^q does not, then consider R_3^c . We have $\{p\} \times \text{pr}_2R \subseteq R_3^c$. Moreover, since by the choice of R relation R_1^q is a non-trivial thick mapping there is $r \in \text{pr}_2R$ such that $(r, c) \notin R_1^q$, hence $(q, r) \notin R_3^c$. Therefore R_3^c is not a thick mapping, a contradiction. Since R_1^q does not equal $\text{pr}_2R \times \text{pr}_3R$ (and $R_2^b \neq \text{pr}_1R \times \text{pr}_3R$, $R_3^c \neq \text{pr}_1R \times \text{pr}_2R$), every η_{ij}^p is non-trivial. Observe that due to the equalities $\text{pr}_{ij}R = \text{pr}_iR \times \text{pr}_jR$, we also have that the unary projections of

R_i^p are equal to those of R for any p ; and therefore all the equivalence relations η_{ji}^p , for a fixed j , are on the same domain, $\text{pr}_j R$. Let

$$\eta_i = \bigvee_{\substack{j \in \{1,2,3\} \setminus \{i\} \\ p \in \text{pr}_j R}} \eta_{ji}^p.$$

Since for any non-crossing α, β we have $\alpha \vee \beta = \alpha \circ \beta$, the relation η_i is pp-definable in R and constant relations without equalities. Since all the η_{ji}^p are non-trivial, η_i is also non-trivial. We set

$$R'(x, y, z) = \exists x', y', z' (R(x', y', z') \wedge \eta_1(x, x') \wedge \eta_2(y, y') \wedge \eta_3(z, z')).$$

Let Q_i^p be defined for R' in the same way as R_i^p for R . Observe that since $(p, q, r) \in R'$ if and only if there is $(p', q', r') \in R$ such that $\langle p, p' \rangle \in \eta_1$, $\langle q, q' \rangle \in \eta_2$, $\langle r, r' \rangle \in \eta_3$, the relations Q_1^p, Q_2^q, Q_3^r for $p \in \text{pr}_1 R', q \in \text{pr}_2 R', r \in \text{pr}_3 R'$ are thick mappings with respect to the equivalence relations η_2, η_3 , relations η_1, η_3 , and relations η_1, η_2 , respectively. All the binary projections of R' equal to the direct product of the corresponding unary projections, while η_1, η_2, η_3 are non-trivial, which means R' is not the direct product of its unary projections, and therefore it is not 2-decomposable. We then can replace R with R' . Thus we have achieved that η_{ij}^p does not depend on the choice of p and i .

Next we show that R can be chosen such that $\text{pr}_1 R = \text{pr}_2 R = \text{pr}_3 R$, $\eta_1 = \eta_2 = \eta_3$, and for each $i \in \{1, 2, 3\}$ there is $r \in \text{pr}_i R$ such that R_i^r is a reflexive relation. If, say, $\text{pr}_1 R \neq \text{pr}_2 R$, or $\eta_1 \neq \eta_2$, or R_3^r is not reflexive for any $r \in \text{pr}_3 R$, consider the following relation

$$R'(x, y, z) = \exists y', z' (R(x, y', z) \wedge R(y, y', z') \wedge C_d(z')).$$

First, observe that $\text{pr}_{ij} R' = \text{pr}_i R' \times \text{pr}_j R'$ for any $i, j \in \{1, 2, 3\}$. Then, for any fixed $r \in \text{pr}_3 R' = \text{pr}_3 R$ the relation $Q_3^r = \{(p, q) \mid (p, q, r) \in R'\}$ is the product $R_3^r \circ (R_3^d)^{-1}$, that is, a non-trivial thick mapping. Thus R' is not 2-decomposable. Furthermore, $\text{pr}_1 R' = \text{pr}_2 R' = \text{pr}_1 R$, for any $r \in \text{pr}_3 R'$ the relation Q_3^r is a thick mapping with respect to η_1, η_1 , and Q_3^d is reflexive. Repeating this procedure for the two remaining pairs of coordinate positions, we obtain a relation R'' with the required properties. Observe that repeating the procedure does not destroy the desired properties where they are already achieved. Replacing R with R'' if necessary, we may assume that R also has all these properties.

Set $B = \text{pr}_1 R = \text{pr}_2 R = \text{pr}_3 R$ and $\eta = \eta_1 = \eta_2 = \eta_3$. Let $p \in B$ be such that R_1^p is reflexive. Let also $q \in B$ be such that $\langle p, q \rangle \notin \eta$. Then $(p, p, p), (p, q, q) \in R$ while $(p, p, q) \notin R$. Choose r such that $(r, p, q) \in R$. Then the restriction of R onto 3-element set $\{p, q, r\}$ is not 2-decomposable. Thus R can be assumed to be a relation on a 3-element set.

If η is not the equality relation, say, $\langle p, r \rangle \in \eta$, then as the restriction of R onto $\{p, q\}$ is still a not 2-decomposable relation, R itself is a relation on the set $\{p, q\}$. Identifying p and q with 0 and 1 it is not hard to see that it is the affine relation $x + y + z = 0$ on $\{p, q\}$. The CSP with global cardinality constraints for this relation is NP-complete by [12].

Suppose that η is the equality relation. Since each of R_1^p, R_1^q, R_1^r is a mapping and $R_1^p \cup R_1^q \cup R_1^r = B^2$, it follows that the three relations are disjoint. As R_1^p is the identity mapping, R_1^q and R_1^r are two cyclic permutations of (the 3-element set) B . Hence either (p, q) or (q, p) belongs to R_1^q . Let it be (p, q) . Restricting R onto $\{p, q\}$ we obtain a relation R' whose projection $\text{pr}_{23} R'$ equals $\{(p, p), (q, q), (p, q)\}$, which is not a thick mapping. A contradiction with the choice of R . \square

REFERENCES

- [1] K.A. Baker and A.F. Pixley. Polynomial interpolation and the chinese remainder theorem. *Mathematische Zeitschrift*, 143:165–174, 1975.
- [2] Libor Barto, Marcin Kozik, and Todd Niven. Graphs, polymorphisms and the complexity of homomorphism problems. In *STOC*, pages 789–796, 2008.
- [3] V.G. Bodnarchuk, L.A. Kaluzhnin, V.N. Kotov, and B.A. Romov. Galois theory for post algebras. i. *Kibernetika*, 3:1–10, 1969.
- [4] Stéphane Bourdais, Philippe Galinier, and Gilles Pesant. Hibiscus: A constraint programming application to staff scheduling in health care. In *CP*, pages 153–167, 2003.
- [5] A.A. Bulatov. Tractable conservative constraint satisfaction problems. In *LICS*, pages 321–330, 2003.
- [6] Andrei A. Bulatov. A dichotomy theorem for constraint satisfaction problems on a 3-element set. *J. ACM*, 53(1):66–120, 2006.
- [7] Andrei A. Bulatov and Víctor Dalmau. A simple algorithm for mal’tsev constraints. *SIAM J. Comput.*, 36(1):16–27, 2006.
- [8] Andrei A. Bulatov and Víctor Dalmau. Towards a dichotomy theorem for the counting constraint satisfaction problem. *Information and Computation*, 205(5):651–678, 2007.
- [9] Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005.
- [10] Jianer Chen and Iyad A. Kanj. Constrained minimum vertex cover in bipartite graphs: complexity and parameterized algorithms. *J. Comput. Syst. Sci.*, 67(4):833–847, 2003.
- [11] M.C. Cooper. An optimal k-consistency algorithm. *Artificial Intelligence*, 41:89–95, 1989.
- [12] Nadia Creignou, Henning Schnoor, and Ilka Schnoor. Non-uniform boolean constraint satisfaction problems with cardinality constraint. In *CSL*, pages 109–123, 2008.
- [13] K. Denecke and S.L. Wismath. *Universal algebra and applications in Theoretical Computer Science*. Chapman and Hall/CRC Press, 2002.
- [14] T. Feder and M.Y. Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory. *SIAM J. Computing*, 28:57–104, 1998.
- [15] D. Geiger. Closed systems of function and predicates. *Pacific Journal of Mathematics*, pages 95–100, 1968.
- [16] G. Gottlob, L. Leone, and F. Scarcello. Hypertree decompositions: A survey. In *MFCS*, volume 2136 of *LNCS*, pages 37–57. Springer-Verlag, 2001.
- [17] Martin Grohe. The complexity of homomorphism and constraint satisfaction problems seen from the other side. *J. ACM*, 54(1), 2007.
- [18] Martin Grohe and Dániel Marx. Constraint solving via fractional edge covers. In *SODA*, pages 289–298, 2006.
- [19] P. Idziak, P. Markovic, R. McKenzie, M. Valeriote, and R. Willard. Tractability and learnability arising from algebras with few subpowers. In *Proceedings of the 22th Annual IEEE Symposium on Logic in Computer Science*. IEEE Computer Society, 2007.
- [20] P.G. Jeavons, D.A. Cohen, and M.C. Cooper. Constraints, consistency and closure. *Artificial Intelligence*, 101(1-2):251–265, 1998.
- [21] P.G. Jeavons, D.A. Cohen, and M. Gyssens. Closure properties of constraints. *J. ACM*, 44:527–548, 1997.
- [22] P.G. Jeavons, D.A. Cohen, and M. Gyssens. How to determine the expressive power of constraints. *Constraints*, 4:113–131, 1999.
- [23] Dániel Marx. Parameterized complexity of constraint satisfaction problems. *Computational Complexity*, 14(2):153–183, 2005. Special issue “Conference on Computational Complexity (CCC) 2004.”.
- [24] Claude-Guy Quimper, Alejandro López-Ortiz, Peter van Beek, and Alexander Golynski. Improved algorithms for the global cardinality constraint. In *CP*, pages 542–556, 2004.
- [25] Jean-Charles Régin and Carla P. Gomes. The cardinality matrix constraint. In *CP*, pages 572–587, 2004.
- [26] I.G. Rosenberg. Multiple-valued hyperstructures. In *Proceedings of the 28th International Symposium on Multiple-Valued Logic (ISMVL ’98)*, pages 326–333, 1998.